

Case Studies 5 – Possibilities and Limitations Moving Forward

Scale, Nuance, and New Expectations in Ethnographic Observation and Sensemaking

ALEXANDRA ZAFIROGLU

Intel Corporation

YEN-NING CHANG

Intel Corporation

We consider new expectations for ethnographic observation and sensemaking in the next 20-25 years, as technology industry ethnographers' work unfolds in the increasing presence of the type of analytical capabilities specially trained (and self-training) machines can do 'better' and 'cheaper' than humans as they can take in, analyze and model digital data at much higher volumes and with an attention to nuance not achievable through human cognition alone. We do so by re-imagining three of our existing ethnographic research projects with the addition of very specific applications of machine learning, computer vision, and Internet of Things sensing and connectivity technologies. We draw speculative conclusions about: (1) how data in-and-of-the world that drives tech innovation will be collected and analyzed, (2) how ethnographers will approach analysis and findings, and (3) how the evidence produced by ethnographers will be evaluated and validated. We argue that these technology capabilities do offer compelling new ways to model and understand the contexts in which ethnographic encounters take place. Yet because ethnography has never been solely about describing behavior, or about testing hypotheses to ultimately generate laws, these new tools will never get us on their own to the type of truths the ethnographer values above all else: the meanings given to experiences by humans.

INTRODUCTION: LIQUEFACTION OF EPISTEMOLOGICAL GROUNDS

Ethnographers employed in various silicon geographies — Valley, Forest, Ally, Wadi, Cape, Gorge, Roundabout¹ — have experienced destabilizing epistemological ground shifts with the emergence of technologies that create new types of data and evidence more valued and trusted by their co-workers to inform product innovation and company business strategies than qualitative, applied ethnography. This epistemological earthquake started with the shudders around big data in the mid aughts, followed quickly by the jostles and rolls of advanced analytics and has culminated in the liquefied ground brought on by the real — and richly imagined *potential* — force of artificial intelligence fueled by data produced by the blanketing of the physical world with the sensing and communication capabilities of the Internet of Things. In geological earthquakes, liquefaction destabilizes the support for building foundations and other objects on the ground; liquefaction of epistemological grounds similarly destabilizes what data count as evidence. In either case, liquefaction is not a permanent state though the ground/data and the objects on it/evidence never revert to their previous arrangements. Instead, attempts to salvage, rebuild, and stabilize on new landscapes ensue. These altered landscapes offer opportunities to refresh or build physical and intellectual structures anew. City planners prioritize structural engineering and update building codes. Technology companies valorize new knowledge systems and data (in the

philosophical sense) for characterizing the world, and new expert workers who create and work with this data (in the computing sense).

In this paper, we explore what will irrecoverably change, what will be contested and what will be the net gains (and losses) for technology industry ethnographers in this new landscape. First, we outline fundamental changes to the narratives the industry tells itself about how innovation happens, and the new competition ethnographers face in observing and making sense of human experiences. We situate this new competition in the larger context of the advent of the fourth industrial revolution and the changes in how data act in the world. Then, through a series of re-imaginings of our current and past ethnography research situated in artificial intelligence and computer vision saturated environments, we will tease out what types of evidence become possible, which types of evidence become more and less relevant or trusted, and re-imagine the role and activities of future technology industry ethnographers.

New Competition in Observing and Making Sense of Human Experiences

In this altered landscape, one certain irreversible change is that silicon geography ethnographers face new competition for their expertise in observing and making sense of human experiences. The established innovation narrative common in the technology sector that justifies the expertise and value of employing ethnographers faces a compelling new narrative in which complex and increasingly ubiquitous computing solutions and systems observe the world, create models of it, make inferences and act with ever-decreasing reliance on direct human intervention. The humans most critical in this new narrative are computer scientists specializing in machine learning and data scientists, experts who can “obtain, scrub, explore, model and interpret data, blending hacking, statistics and machine learning” (Woods 2011²). In 2010, Hal Varian described statisticians as holders of “the sexy job in the next ten years” (McKinsey 2009). By 2012, *statistician* had morphed into *data scientist*, in recognition of the computer science and machine learning expertise entailed, and sexy morphed to sexiest with a much-extended expiration date in a widely cited Harvard Business Review article provocatively titled “The Sexiest Job of the 21st Century” (Davenport and Patil 2012). With this new innovation narrative sweeping the technology industry and beyond, LinkedIn reported machine learning engineer and data scientist lead the fastest growing job roles in the United States in 2017, followed closely by big data developer (#5) and director of data science (#8) (Bowley 2017). In early 2018, GlassDoor announced data scientist the “top job” in the US for the third year in a row, based on the number of job openings listed on their service, the salary levels reported by those with this job title, and these workers’ overall satisfaction with their jobs (Glassdoor 2018). As the liquefied ground for what data count as evidence for innovation re-stabilizes, innovation-focused ethnographers working in/on/around/about the technology industry (and others) have reacted how we might expect from aggrieved qualitative social scientists; with denial (qualitative data matters!), with bargaining (we need both big data and thick data!); with testing (how do we combine quant and qual data in new ways?) and acceptance (it appears to be here to stay, so let’s do what social scientists do best and cut it down to size by interrogating our assumptions, our language and what’s actually possible now)³. Here, we add to these voices with a contribution to the far end of the acceptance literature, as we imagine the new concretized

landscape and how silicon geography ethnographers will work there in the near and far futures.

We refer to the established technology innovation narrative as *the people-centric narrative* and the new one as *the data-centric narrative*. These are the **idealized stories** the industry tells itself and others about how it uses evidence in-and-of the world to create new products and services customers will value and pay for. As idealized stories, they are neither nuanced nor subtle and thus not true to how teams of diversely trained researchers often collaborate at technology companies, as we will illustrate later. While they grossly oversimplify how work actually happens, they do provide a useful framework to highlight differences in the underlying idealized *logic, methods, analysis, outputs* and *end goal* flow for innovation processes in the technology sector. The coexistence of these differences constitute the unstable ground silicon geography ethnographers currently find themselves occupying.

The people-centric narrative emphasizes the **logics** of opportunity areas for technology intervention and the existence of unmet needs, **methods** for their discovery through qualitative, ethnographic inquiry and **analysis** informed by social science theory with an **output** of implications for design then acted upon by product and strategy teams⁴. It is an unrelentingly human-centric model. The unmet needs may be of atomized individuals (consumers, workers) or of social and/or economic organizations (households, enterprise) but the **end goal** is always addressing the desires and needs of people. This narrative now co-exists, often in epistemological friction, with the new data-centric narrative of innovation. This narrative privileges the **logic** of powerful yet invisible patterns all around us, **methods** for their discovery through instrumenting the world to digitize the stubbornly analog and sucking up massive amounts of resulting data which are then combined with existing digital data. These data sets are then cleaned and readied for **analysis** by data scientists and computer scientists using new, increasingly complex and progressively more opaque algorithmic processes that then **output** models of the world that machines can use to directly understand and act in the world. The **end goal** in this narrative is to create compute systems that exhibit a specific behavior: artificial intelligence (AI). This innovation model is unrelentingly machine and data centric, with the people so important in the earlier narrative relegated to the edges; as (often unwitting) providers of behavioral inputs to feed the voracious data appetites of increasingly autonomous intelligent systems, and as subjects acted upon by such systems as they go about their daily activities — usually polarized between delighted (in corporate vision videos) or alarmed (in critiques of this narrative) (see O’Neil 2016).

This data-centric innovation narrative is part and parcel of the broader narrative about the fourth industrial revolution and confluence of big data, advanced analytics, machine learning and its subfield of computer vision, the Internet of Things (IoT) and other technologies that in combination enable machines to exhibit artificial intelligence, and that increasingly blur the lines between the physical, the digital and the biological (Schwab 2017). Various potential futures await earth’s inhabitants, and we are challenged to decipher which voices in the popular debates around AI are Pollyannas, Chicken Littles, or Cassandras, and which are simply exploiting the current context in which “technologists, businesspeople, and journalists wield the idea like a magic wand that turns ordinary computer software and devices into world-saving (or world-ending) marvels” (Bogost 2017), simply to further their personal or company brand strategies.

The envisioned full expression of the fourth industrial revolution is the arrival of High Level Machine Intelligence (HLMI), “achieved when unaided machines can accomplish every task better and more cheaply than human workers” (Grace et al 2018), a clever definition as it is broad enough to be both meaningless (what’s a task? how is cost measured?) and to genuinely alarm observers. A recent survey of machine learning experts that explored their beliefs about how quickly the world is progressing towards HLMI suggests (if the experts are right) that we have roughly 120 more years before AI can automate all human jobs. These surveyed experts expect AI systems to be able to tackle more and more ambitious analytical and creative tasks in the intervening century, such as folding laundry “as well as and as fast as the median human clothing store employee” (in 5.6 years), assembling any Lego (in 8.4 years), generating a top 40 Pop song (in 11.4 years) or writing a New York Times bestseller (in 33 years) (Grace et al 2018). Considering the work done by the average silicon geography ethnographer to be somewhere between writing a pop song and writing a best seller, we may have approximately another 20-25 years of work ahead of us before the machines takes over.

Of course, we are being facetious. We believe the future that technology industry ethnographers need to consider is not when and how sentient machines will take our jobs, enslave, or murder us. Rather, we need to consider new expectations for ethnographic observation and sensemaking in the next 20-25 years, as our work unfolds in the increasing presence of the type of analytical capabilities specially trained (and self-training) machines can do ‘better’ and ‘cheaper’ than humans as they can take in, analyze and model digital data at much higher volumes and with an attention to nuance not achievable through human cognition alone. As we examine these new expectations here, we will avoid clichéd tropes about the nature of work, Pollyanna, Chicken Little or Cassandra proclamations about artificial intelligence, as well as any mention of sentient robots. We will contemplate these new expectations by imagining and then dissecting very specific applications of machine learning, computer vision, and Internet of Things sensing and connectivity technologies to our current and past ethnographic projects. We will use these case studies to draw speculative conclusions about: (1) how data in-and-of-the world that drives tech innovation will be collected and analyzed, (2) how ethnographers will approach analysis and findings, and (3) how the evidence produced by ethnographers will be evaluated and validated. These imaginings and dissections will highlight the current frictions in approaches to making sense of the world between technology industry ethnographers and some of their workplace collaborators, stakeholders or audiences, namely: engineering, computer sciences and machine learning experts, as well as marketing teams and business leaders. Given their bases in long-standing disciplinary differences and the hype around big data and the ‘magic wand’ of AI, these epistemic frictions are unlikely to be resolved. However, as is the case with most frictions, we believe these epistemic ones will be productive. So before we turn to our imaginings, we will first consider three key differences in how we as technology industry ethnographers and our workplace collaborators approach observation and sensemaking of the world. These differences provide the context for how we imagine working in the world saturated with the types of sense and sensemaking technologies that our colleagues are creating.

EXISTING FRICTIONS AND THE RUMBLING INNOVATION LANDSCAPE

As technology industry ethnographers, we have endured repeated and sometimes arched-eyebrow questioning by our colleagues about how we observe and make sense of the world. We have had (our own!) sales people introduce us to customers as appetizers (tasty and insubstantial) before the main course (meaty and serious) of technology discussions. We've been informed that our work is fluffy (not important), and our data anecdotal (not evidence). We are familiar with the polite thanks for sharing our interesting (not useful) work at the conclusion of meetings. We've been told that what we do is simply descriptive observation (no theory needed), for which advanced training is not truly necessary and that we should/can be easily replaced with cheap interns. The common misconception behind these critiques is an assumption that we are pure empiricists dedicated solely to inductive reasoning — that we take a highly limited amount of data and reason directly from it to grand conclusions. The theory that informs our data collection and analysis strategies is unrecognized by our peers primarily for two reasons. First, because we eschew name dropping theorists and alienating stakeholders with inside baseball explanations of the theoretical foundations of our research decisions as a starting point for discussions of research findings⁵. Second, because the blend of abductive, deductive and inductive reasoning we favor differs substantially from established methods of observing and making sense of the world rooted in models of stating, testing, and rejecting or accepting hypotheses, and to the types of empirical data favored in the engineering and computer science disciplinary training of our colleagues. After recently presenting a thoughtfully considered multi-pronged research program incorporating very specific proposed research projects to identify near, mid and long term strategic and tactical opportunities for a business line, we were criticized for not understanding how research works (despite our Ph.Ds.). A primary stakeholder corrected our approach, offering that in order to do research we needed to identify hypotheses to test, and proceeded to give us several examples of how to correctly formulate a testable hypothesis.

In full disclosure, these occasionally disheartening — but mostly amusing — encounters are generally with internal stakeholders removed from day-to-day R & D work, and not the computer scientists and engineers we work with closely and for extended periods of time on innovation projects. Turning to these latter colleagues, we recognize some fundamental differences in the assumptions we bring to creating and using data in and of the world to inform innovation. The frictions our differing assumptions create has been overwhelmingly productive — getting us to new questions and new knowledge. We will use examples from one project, *Local Experiences of Automobility* (LEAM for short) that we will revisit later as one of our future case studies, to illustrate these frictions. The LEAM team, led by one of the authors, consisted of anthropologists, an electrical engineer, a psychologist, computer scientists, and a data visualization intern, and is an excellent example of close collaboration between colleagues with different disciplinary training and of the epistemological frictions that ensued and that will likely be exacerbated in coming years as the ground concretizes around data-centric innovation.

You Are Giving Us Too Many Problems

When we started LEAM in 2010, our idea of a suitable ethnographic research question for exploring the future of autonomous vehicles was simple, open-ended, and broad: *what is a car?* This left us open to a number of research methods and tools that would generate many types of data to identify what we called opportunity areas for the development and application of new technology solutions for road-based transportation. As social scientists we assumed our job was to generate data to identify *opportunities* to make change in the world. As computer scientists and engineers, our colleagues assumed their job was to generate and use data to *create* change in the world. One teammate in particular was fond of reminding us:

Engineers work in the solution space. We like to solve problems. Anthropologists like to find new problems, even when we are still working to solve the ones you gave us last week. You are giving us too many problems.⁶

Her idea of an appropriate question was: *40% of the build of materials cost of a car is sensors and electronics. What if the sensors were so good that a car couldn't be hit? Then it could be an egg – no airbags, no crash frame, and no bumper.* This approach enabled her to focus specifically on increasing the value derived from ingredients currently going into a solution — solving the problem of justifying the BOM (build of materials) cost of a car devoted to sensors, by decreasing the amount devoted to bending metal. As a team, we overcame this friction by learning to announce at the beginning of activities which mode we were in: problem or solution, and by prioritizing the problems we identified so as not to overwhelm our team.

This guy can't be trusted! This data is no good. Throw it out!

As the lineup of researchers was still forming in the earliest days of LEAM, we had team members each review a single ethnographic interview transcript with a car owner as a means of building empathy and awareness of current automobility practices. Upon meeting to talk about what we had read, one of our teammates, a computer scientist highly skilled in machine learning, quickly pointed out that the participant in his transcript contradicted himself in describing how he uses his car. He then dramatically threw his hands in the air and declared: “This guy can't be trusted! This data is no good! Throw it out!” For him, the data was bad; how could he build a model on logically inconsistent data? In both ethnographic analysis and in building machine learning solutions, data needs to be assessed, cleaned and criteria established to determine what data to include and what not to include in the analysis. In this case, our computer scientist teammate applied the rules of his training to a data type that didn't follow the validity criteria needed to build a predictive model. Our inability to convince him that such inconsistencies did not invalidate the data, they merely suggested a different type of analysis than he was accustomed to, meant his stint as on the team ended there.

Excellent! We'll get to ground truth

When we started scoping the research methods for LEAM, we recognized that because we were after the mundane, taken-for-granted details of owning, using, caring for and being around private passenger cars, interviews alone were insufficient. We settled on a mix of

methods that we have described in detail elsewhere (Zafiroglu, Healey, Plowman 2012), and that we will briefly summarize here as semi-structured interviews and ‘carchaeology’ inventories of objects in and on vehicles on the one hand (getting us to ‘Remembered Drives’) and 30-day collections of GPS data, mobile phone use during and around drives, and car-interior sensing of sound pressure, lighting, accelerometer, temperature (getting us to ‘Recorded Drives’).

Feeling confident that these methods would allow us to create thick data from two thin sources; the after-the-fact remembrances and explanations provided in interviews and during carchaeology sessions, and the machine-produced numbers and patterns from the sensors, GPS and phone monitoring app, we were taken by surprise by the interpretation of our methods by our electrical engineer team mate. Referring to the Recorded Drives data set, she proclaimed her excitement that we would finally get to ‘ground truth’ about how people were using their cars and their phones while in their cars. Concerned (like our erstwhile computer scientist team member) that people are inconsistent and lie when they talk to us, she argued the machine data would provide better evidence for how people *actually* used their phones. We could, in fact, catch participants when they lie! This was a goal the ethnographers hadn’t even imagined. What ensued were months of animated discussions about ‘truth’ in research data and how we create and trust data created through our interactions with people vs. people’s interactions with machines. Our colleague’s confidence in the empirical objectivity of digital traces created by people through their interactions with technology to get us closer to what we were trying to measure — to ground truth — was never truly reconciled with the social science trained ethnographers keen interest in the multiple truths people produce and live in simultaneously. A data point from a participant in China drove home our different takes on evidence. Our GPS and phone monitoring data showed the participant using the camera function on his mobile phone during his commute to work while travelling at relatively high speeds. When we queried him about taking pictures while driving, he explained he was merely placing his phone in a holder attached to his dashboard and as the camera button was on the side of the phone, he sometimes inadvertently took a picture. For our computer science and engineering teammates, the data point was errant; while it did get us closer to the ground truth of when a phone was handled while driving, it was incorrectly labeled by our phone monitoring software and therefore as currently labeled would not be useful for training a machine learning model. For the ethnographers, the phone monitoring data point wasn’t errant; it was simply data about a mundane detail of phone and car use the participant would not have thought to mention to us and that we would not have known to ask about without the machine data. Of these three frictions, this was the most difficult for us to reconcile as a team, as half us wanted ‘ground truth’ data useful for training a machine learning algorithm and the other half was just being exposed to the idea of machine learning and the data rules it requires.

Past the Shudders, Jostles and Rolls and Onward to Liquefaction

In addition to these long standing frictions between social science and computer science and engineering research methodologies and understandings of evidence, in recent years we’ve lived through the transformation of our company, like so many others, to being data-driven. ‘Data-driven’ attests not to the type of evidence produced by qualitative ethnographic research, but by big data, advanced analytics and IoT capabilities

that are creating the massive amounts of data that drive the new data-centric innovation model. In a recent keynote for a “Data-Centric Innovation Summit” hosted by our company, an executive vice president presented a “vision for a new era of data-centric computing” and explained the opportunity before us as the biggest TAM (total available market) growth opportunity in our company’s history.

I think one of the most stunning statistics is that 90% of the world’s data has been created in the last two years and even more stunning perhaps is that only about 1% of that data is being utilized to create any sort of meaningful business value. I see tons of room for our industry to grow. (Intel 2018a)

This VP’s enthusiasm succinctly illustrates the ascendancy of empiricist models for innovation fueled by the increasing generation and availability of big data in business. This is a compelling model, if as Kitchin argues, those in business believe “the volume of data, accompanied by techniques that can reveal their inherent truth, enables data to speak for themselves free of theory” (2014, 3). With the adoption of the data-centric innovation narrative, the evidence needed to fuel innovation has shifted from being *about people and their lived experiences* to being *digital traces created by people*, often at very high scales and from multiple sources. The outcome, as Metcalf and Crawford argue, is “the familiar human subject is largely invisible or irrelevant to data science” (2016, 3). The vagaries and dissimulations of the familiar human subject no longer matter if the focus is on the (imagined) veracity of digital traces. If researchers focus on *digital traces* that are imaged to be objective and capable of speaking for themselves then, (for example) a contracting eyelid is empirically a contracting eyelid that can be analyzed using new data analytics methods; there is no need to engage with the whole of the person attached to said lid. In the age of computer vision and big data Geertz’s argument for thick description becomes particularly troubling for the ethnographer.

Consider . . . two boys rapidly contracting the eyelids of their right eyes. In one, this is an involuntary twitch; in the other, a conspiratorial signal to a friend. The two movements are, as movements, identical; from an I-am-a-camera, “phenomenalistic” observation of them alone, one could not tell which was a twitch and which was wink, or indeed whether both or either was twitch or wink. Yet the difference, however unphotographable, between a twitch and a wink is vast; as anyone unfortunate enough to have had the first taken for the second knows. (Geertz 1973, 6)

In the 45 years since Geertz wrote about “I-am-camera” observation, the world has profoundly changed; such differences may no longer be “unphotographable”. We can now imagine a world in which a camera connected to computer vision capabilities can distinguish a wink from a blink, if trained with a sufficiently large data set of humans contracting the eyelids of their right eyes. While conceivable, we are still left wondering *why*, both in terms of the significance of the contraction (a sign of disease? a sign of playfulness? serious ill intent?) and in terms of what the return on investment might be for developing and training such a model (who would pay for it and why?). In other words, as Kitchin notes, “while data can be interpreted free of context and domain-specific expertise, such an epistemological interpretation is likely to be anemic or unhelpful as it lacks embedding in wider debate and knowledge.”(Kitchin 2014, 5)

And yet, in our work environment as technology industry ethnographers, there exists tremendous confidence in AI systems built on computer vision and internet of things sensing that promise to model and interpret the cyber-physical world and to automate decision making and actions with limited human contextual and domain-specific expertise needed beyond the initial set up of the systems by data and computer scientists. Big data, advanced analytics, AI and IoT promise to simplify and speed up innovation, by creating a “new mode of science, one in which the modus operandi is purely inductive in nature” (Kitchin 2014: 4). Our colleagues believe this and our leaders espouse the primacy of data and such solutions to our future financial success.

Regardless of our clear understanding of why such a purely inductive science is preposterous⁷, as ethnographers our work will be increasingly situated and carried out in contexts where AI, computer vision and machine learning algorithms are constantly sensing and building models of people, activities and objects, that feed services which will significantly reconfigure our social, political, and economic activities and our daily behaviors and interactions. In the not-so-distant future, we will work, move through, live, shop, recreate and more in environments where machines will create models independent of human interpretation to infer what’s happened, what’s happening, and what’s *likely* about to happen, or to *create* what happens by arranging local conditions towards outcomes desired by those deploying such tools. As Paglen argues, those deploying computer vision and related tools will be able to “exercise power on dramatically larger and smaller scales than has ever been possible”, seemingly objectively as the ideological foundations of the algorithms informing interpretations of images “function on an invisible plane and are not dependent on a human seeing-subject”(Paglen, 2016). 2017 news coverage of toilet paper dispenser kiosks using facial recognition to limit patrons to nine sheets of tissue each per 15 minutes at Beijing’s Temple of Heaven is an excellent example of one such exercise of power.

As ethnographers employed in the technology industry, we face a challenge and responsibility as we may both be *acted upon* by such systems and our work (ideally) will *shape* how these systems and data exercise power by contributing to how they are created, managed, updated, and scrapped. Considering how we may be *acted upon*, the presence of these systems in field sites and in our work environments means we face existential and tactical questions about our practice. The growing literature on the existential questions addressed pressing concerns including: are our skills still needed and valued for innovation (Madsbjerg 2017)? Can our employers replace us outright with algorithms and AI systems and a few data scientists (until the machines become ‘good enough’)? Does the technology industry need fewer ethnographers and more human factors engineers who study how experts will build such systems? Or who study how their business customers will interact with the outputs of these systems? What are the burning questions that will keep ethnographers relevant and employed (boyd and Crawford 2012)?

In the remainder of this paper we will focus on tactical questions we face as ethnographers being acted upon by the technologies of the fourth industrial revolution, including: How will AI and machine learning applications such as computer vision, text analytics, and speech understanding reshape how we collect and analyze data in-and-of-the world that drives tech innovation? What data will and will not count as evidence? How will the evidence produced by ethnographers be evaluated and validated by our colleagues and by our employers potentially using these self-same technologies and tools? In short, how will

expectations for ethnographic practice in the technology industry change in the next 20-25 years?

ETHNOGRAPHY ON THE (NEW CONCRETIZED) GROUND

To answer these questions requires a bit of imagination, but not a full flight of fancy into Chicken Little proclamations about sentient job-killing robots that we promised to avoid. Our imagination of the future needs grounding in our current ethnographic practices and the technology capabilities at our disposal today and in the foreseeable future. Therefore, we will cut down the abstract and idealized fourth industrial revolution and AI into more realistic yet still ambitious applications of machine learning, computer vision, natural language and audio processing, and Internet of Things sensing and connectivity technologies that we then imagine included in three real (not idealized) ethnography projects we have undertaken in the tech industry.

We will avoid declarations about the exact timing of the blanketing of the earth in these new sense and sensemaking technologies, and instead imagine each project in a *low*, or a *medium* or a *high* presence of advanced sensing and sensemaking capabilities present in our research settings and in our work settings (i.e. where we analyze our data).

For each case study, we will describe the original project, our methods, the data we produced and the project outcomes. We will then re-imagine the project with new combinations of sensing and sensemaking capabilities in our research and data analysis locations and explore how adding these would have first order and second order consequences. Here, first order consequences refer to what will change in ethnographers' research practices and methods; how we will approach sense and sensemaking in our research and how data will and be created to achieve our research goals. In contrast, we define second order consequences here as how professional expectations for technology industry ethnographers will change; what may be new expectations for future ethnographers' responsibilities and scope of work, and new professional standards for training and for how our work is evaluated and validated. We will avoid long discussions on the mechanics of our access to such new data. We will assume that access to research participants' own data will be arranged with and consented to directly by them, and access to some subsets and/or versions of broader public or private data sets (such as security camera data from a housing complex; utility usage data for a neighborhood) will be arranged with the data set owners, and the uses to which we put these data sets will follow the ethical guidelines of the American Anthropological Association, and the privacy requirements and research approval process of our employer, which jointly define our current research practice.

Table 1. Imagined Levels of Advanced Sense and Sense-Making Capabilities

Presence of Advanced Sense and Sensemaking Capabilities	In Field Locations: Data Creation, Data Access and Sense Making Capabilities	At Ethnographers' Office Locations: Tools for Sensemaking	Interaction Between Computing Systems and Ethnographer
Low	Access to historic data such as utility usage, home access systems as provided by participants or by service providers	Access to computer vision, big data analytics, speech and audio detection and interpretation systems that we use to analyze the field data we have already created, our participants have already created, and stored data already generated by others (utility usage, for example)	Computing systems alert ethnographers to patterns or anomalies within field data, and in field data within context of larger data sets
Medium	Close to real-time access to existing data being collected and analyzed by service providers and other actors. Ethnographers' field equipment has machine learning and computer vision capabilities that can act on data as it is created and present analysis to researchers.	Same as low level	Computing systems proactively suggest topics and follow up questions to ethnographers based on patterns and anomalies in field data.
High	New sensors, devices, connectivity and computing solutions that create models of environment and can respond with real-world action to local conditions and to human or machine commands	Access to advanced data analysis and research management systems	Computing systems generate protocols and questions based on priorities and inputs from ethnographers, and present for approval; computing systems automate some interactions with research participants

Smart Home Economics (SHE) and 21st Century Homemaking

In 2014, our strategists and product development teams considered home security the entry point for the smart home market, from which we could add on home automation capabilities such as automating locks, lighting, and HVAC systems. After security and basic home automation, what other usages would be valued? As ethnographers, we set out to think beyond security and simple automation to other experiences that could be possible, appropriate, and valued in homes equipped with new IoT sensing and sensemaking through field visits with householders and with home services professionals. In Smart Home Economics (SHE) we explored possibilities around practices of caring for the physical structure of a home and around managing a household.

Our field methods reflected our goal of thinking beyond incremental improvements to existing home security products. We explored *homemaking* activities at two scales. At a large scale: how are homes purchased/leased/otherwise come to be occupied, finished, maintained and updated? How do people come to live where they do? What processes are

involved? At a small scale: what are considered normal and necessary daily practices in homes and how are they achieved through people, devices, systems, expectations and conventions? We visited eight households each in greater Shanghai and greater Atlanta. We spent an initial three hours interviewing in homes, then had participants answer personalized follow up questions via a smart phone app with which they took videos and pictures and answered questions. We then revisited them a week later and reviewed their follow up answers. These interviews were in English or a combination of English and Mandarin (using a translator) and were video recorded. In parallel, we conducted in-person interviews with home services professionals about their work practices, skills, perspectives on what was changing in their markets, and their outlooks on how digital technologies were impacting their work. These experts included: real estate agents, property managers for luxury, non-luxury apartment, and other housing complexes, home inspectors, building and home automation solution designers and installers, home and building security specialists and installers, security guards for apartment complexes, construction managers, interior designers, and architects. These interviews were in English or Mandarin. Atlanta expert interviews were video recorded and Shanghai expert interviews audio-recorded.

SHE resulted in a data-informed critique of then current home automation and security products. We contended such products were chasing the tail end of the 20th century by seeking to further digitize homemaking practices that had already been almost fully automated in the past two centuries. Have a vacuum? Make it robotic! Have a dishwasher or washing machine? Connect it to the internet! Have electric lights? Control them from afar! We argued to truly think beyond relatively small tweaks to previously achieved dramatic improvements in homemaking practices, 21st century smart home solutions must position householders for domesticity in a world with very different environmental, political, social and economic contexts than the previous century. We offered experiential statements for domestic life in these new contexts and different priorities around security or automation usages responsive to new concerns and situated in broader networks of services, systems and actors within and beyond the home. We followed SHE with SHIFT (Smart Home Information Flow Technologies), an ethnographic project on householders' expectations and preferences for sharing information about daily homemaking activities with others⁸. SHE and SHIFT were central to the comprehensive Smart Home usage roadmap we authored that our business and technology teams relied on extensively when defining new product capabilities for the smart home market.

SHE in a Low Sense and Sensemaking World

Let's now reimagine SHE with a low presence of advanced sensing and sense making technologies in the homes and work locations we visited, and in our own work settings as we analyze the data. In this alternative reality, we imagine 'low' presence in *our field locations* simply means that we have some level of access to stored (not real time) data already being generated by other actors that we didn't have in 2014, such as:

- video and audio files from apartment complexes, residential neighborhoods and individual home security systems.

- event logs from security, home access and communication systems, such as when key passes were used to enter a gated community, or park a car in a garage, or when calls were made to an apartment from an entry call box.
- Current and historic logs of utility usage in homes; patterns of electricity, heating, gas, water consumption as captured by service providers for our participants and for the areas they live in (an apartment complex; a city; a region)

We imagine as we analyze the research data in *our own work settings*, 'low' presence means we have access to computer vision, big data analytics, speech and audio detection and interpretation systems that we could use to analyze the field data we have already created in person (our interview video and audio recordings, our still images, our audio field notes), our participants have already created (video and text data created through the smart phone app) and the historic utility use, security footage and event log data from home settings we now have access to. In other words as the level of sense and sensemaking technologies are low, we assume we will *not* have real-time access to any conditions, actions or events happening in our field locations, we will simply be able to analyze our existing data differently after it has been digitally collected and stored.

So, what might meaningfully change in terms of first order consequences, i.e. how we conduct research? Given that we will not have real time access to machine data, the main differences in how we sense and sensemake will happen as we review and analyze existing data types in office. Currently, the audio and visual recordings we create and our participants create on our behalf are primarily useful as a way to supplement and extend *our human powers of observation and memory* of events, activities, interactions, locations, objects and actors. We review images and videos and audio transcripts in order to recall details we may have failed to originally observe or simply forgotten between visits and analysis days or weeks later. How was that living room arranged? What exactly did that broken refrigerator ice maker work around look like? Were there informative turns of phrase we didn't notice during an interview? An off-hand comment that we didn't quite grasp or didn't realize was noteworthy (pun intended) at the time but on review significantly informs our analysis? If we were able to analyze this data using computing systems trained to detect patterns or anomalies, these digital artifacts would shift from simple observation and memory aids *for us* to potentially powerful tools to observe at nuances we can't detect and to sensemake using massive memory of other events we do not personally possess. With trained models, we could open our field data to new machine interrogation which might be able to detect patterns or significant events we captured but failed to notice. In other words, the trained machines may be able to see what we couldn't in our videos, or hear what we couldn't in our recordings. In the majority of current applications of computer vision analysis of video in law enforcement and commercial security services, such use is forensic; seeking after-the-fact evidence of a crime committed. Here we might use forensic in a different way to call attention to the 'crime' of leaving data collected by ethnographers unnoticed or unanalyzed. We imagine the following five ways we as ethnographers may forensically employ AI systems in office to interrogate our field data.

We See In Office What We Couldn't See While In Field – Using computer vision enabled computing systems in our offices, we could be alerted to objects, movements, behaviors, activities that we were not attuned to in field and still do not or cannot notice later when we

review images and videos. We imagine our office systems producing reports on field work video footage and still images that provide descriptions or categorizations of people, objects, and conditions using criteria not available to the unaided human eye alone. Concerning people, this could include such things as patterns in facial expressions (perhaps glossed as ‘emotions’), body posture, and the body language of participants and of researchers. For objects, this could include recognizing objects or arrangements of objects in the environment that are common or uncommon to a larger household demographic, or to create novel classifications of household types based on visual evidence of homemaking practices that we might not have reached on our own (or not reached as quickly). For environmental conditions we cannot visually analyze with precision on our own, this could include evaluations of indoor and outdoor air quality based on our video or still images (see Zhang et al 2016) or standardized evaluations of patterns in lighting practices and conditions in homes using metrics valued by our engineering colleagues. We could achieve the attention to the physical environment Reichenbach and Wesolkowska (2008) argue is often missing in ethnographic research.

We Hear In Office What We Were Not Attuned To While In Field – We imagine running our video and audio through machine learning audio software that analyzes non-speech sounds that we didn’t notice, that we tuned out over the course of fieldwork, or that were inadvertently captured on video recordings produced for us by participants. These could include inside-the-home and outside-the-home soundscapes of domestic life, from appliances or consumer electronics running, to heating and cooling systems turning on and off, to neighbors moving about, to traffic, garbage pick up, deliveries, or grandmas blasting music while practicing Tai Chi in the condo complex courtyard. Such software could listen and classify these sounds for us, and are not so fanciful given current product offerings from companies such as Audio Analytic, SoundHound and several others. Reports on these sounds could turn unrecognized sounds into recognized ones, and could spark us to ask new questions about homemaking. While beyond the capabilities of current audio AI product offerings, we know university researchers have shown how computer vision analysis of small movements of inanimate objects (a potted plant, an empty potato chip bag) in a silent video can be independently used to recreate the audio (including speech) occurring when the video was created (Feltman 2014). We imagine running existing video footage from security cameras, which often don’t include audio, through computer vision algorithms to identify audio events and patterns around home exteriors — from multi-tenant dwellings to single family homes — that could spur us to ask more informed questions about home security conditions and practices.

We Understand Speech That We Didn’t Understand In Field – We imagine running our audio (or video) files through AI language translation and analysis software that can flag points in conversations where misunderstandings may have occurred, so we can decide to explore more with participants in the follow up interviews. Such language translation and analysis capabilities would be useful between completely different languages (Mandarin and English) and between dialects and slang within a language (Atlanta/SE United States English and West Coast English). We may be alerted to regional accents we don’t recognize, or to idioms we don’t recognize as having local significance. We can also imagine needing to train the software to recognize intentional miscommunications; when a topic is skirted to save face, or

because it is too personal, or a follow up question is not needed as the people present already recognized and sorted the miscommunication non-verbally.

We Contextualize Our Data In Ways Satisfactory To Our Colleagues – The memory used by machine learning systems (the sets of data on which a machine learning algorithm are trained) is much larger than we as humans can retain and recall. Given we will have access to home utility data, as well as audio/video we have captured, we imagine that our machine learning tools can help us understand if a research participant is representative of a larger group of householders. Is the amount and pattern of utility use of this household typical or unusual for their neighborhood or complex? Or an extreme in some way? Using computer vision capabilities, could we understand if there is an object present, a home decorating style, a pattern in the arrangement of objects, or the range of objects in this home that have social, cultural, economic, or political significance in the field locale? As researchers, we value both participants that represent larger groups as well as extremes in behavior, product ownership, income, etc. The machine learning software will help us better contextualize who we have included in our study.

We Extend The Usefulness of Ethnographic Data Over Time – In addition to subjecting the data from SHE to analysis by machines, we may consider adding the SHE data to the broader data set that feeds the machine learning models we've used in the study. This of course raises a number of questions about how we do this, for which we will need to collaborate closely with data scientists. Beyond the obvious difficulties of cleaning, structuring and properly labeling the data, we face practical questions about data handling. How do we write a consent form for participants that encompasses use in training models? How do we explain who has access to the raw data in the future? Currently, we specify “only the immediate research team”. Will we need to ask for consent for additional researchers? For machines? What might that look like? Furthermore, several years later, would we want to — and legally and ethically could we — re-run the data from SHE through updated machine learning tools to see if other insights are possible with retrained and updated AI tools?

These new in-office analysis capabilities will create novel data artifacts in the form of text-based reports that we will treat similarly to our existing field notes, field audio transcripts, and image data; as another input to be analyzed. They won't replace the type of social theory informed analysis that we current undertake, but they will change two aspects of our research practice. First, we will extend our ability to identify areas of interest based on nuance and scale of observation that we did not have before. Second, we will create richer, more scientifically precise descriptions of the locations for which we are designing new technologies that we can share with product development teams.

Given these outcomes, what might meaningfully change in terms of second order consequences, i.e. professional expectations for technology industry ethnographers? We see three potential outcomes.

Ethnographers Face New Expectations For Proving Data Validity – We will now have better means to explain who our research participants are. With analysis of our participants demographics and behaviors in comparison to larger data sets we can get closer to ground

truth with our internal engineering trained research audiences about how representative, or how unique, our participant sample is. We can now better explain and support our decisions for who we have included in our studies, and our data are less likely to be doubted as anecdotal by our colleagues.

Ethnographers Face New Expectations To Wring More Insight From The Same Type and Amount of Qualitative Data – Because we can now perform forensic data analysis on our standard digital research files that extends the scale at which we can analyze our data and the nuance in observation beyond our human sense and sense-making abilities, the type and breadth of deliverables from a single project will likely increase. As a simple example, the types of data amenable to creating ‘ground truth’ to train algorithms by our computer science colleagues will be expected as part of research findings. (how well-lighted are American vs. Chinese homes?)

New Professional Standards For Ethnographers’ Skills And Fluency In Engaging and Contesting Machine Data – Ethnographers will be expected to be able to parse the types of reports produced by machine learning systems. Even if the front end UI available for accessing the analysis on our work machines resembles those used by consumer wearable or smart home services, ethnographers will need to be able to look behind the UI and be able to understand enough of how the analysis was performed to contest or at least understand how the machines came to a conclusion. This will particularly be an important skill if the ethnographer independently reaches a different conclusion than the trained machine, and the ethnographer needs to give input on how the training models should be updated.

Local Experiences of Automobility (LEAM) and the Future of Transportation

Imagining a bit more intense presence of sense and sense-making technologies in our field and work environments we return to Local Experience of Automobility (LEAM) from which we drew some of our earlier examples of disciplinary frictions. The ultimate goal of the research was to prepare Intel to design vehicle and transportation system solutions as we entered a decade of transformation of cars, road infrastructure and ecosystems through advanced sensing technologies, computational systems and services. While we did several rounds of research in seven countries, the richest methodology was used with car owners in two cities each in Brazil, China and Germany. With these participants, we completed car inventories, semi-structured interviews, video diaries and 30 days of car use data including GPS data, car-interior sensing of sound pressure, lighting accelerometer, air temperature and use of mobile phones before, during and after drives. Participants were visited once for an initial three hour interview, car inventory and ride-along, and to have sensors and GPS installed in their car, and tracking software installed on their phones for thirty days. They were revisited again approximately 6-8 week later to review the phone and GPS data the research team had visualized using Google Maps.

The outcomes of LEAM included the generation of over three hundred use cases for in-vehicle infotainment systems, advanced driver assistance systems and semi-autonomous driving solutions. The project was notable because it created a direct tie between foundational qualitative research and product definition, as well as generating over forty awarded patents, and multiple internal prototypes and projects with car manufacturers.

LEAM in Medium Sense and Sensemaking World

Let's now imagine LEAM with a medium presence of advanced sensing and sense making technologies in our research equipment, in the cars and the road systems and transportation infrastructure in the six cities we visited in China, Germany, and Brazil, and in our office settings as we analyzed the ethnographic data. In this alternative reality, medium presence in *our field locations* might mean that in addition to the 'low' presence capabilities of the last case study, we have close to real-time access to existing data that is being collected and analyzed by other actors such as:

- Current road conditions and traffic patterns
- City-level mobile phone location and use data, including phones being used in moving and still vehicles
- Histories of a research participants' cars presence on and use of roads during the past 6 months (based on phone data or transponders for road fees), presented on digital maps, and anonymized comparisons to other car owners and averages in the municipal area

And it might mean that the cameras and audio recording systems we bring with us to the field incorporate computer vision and machine learning capabilities that act on or analyze data as we create it and present analysis to us by showing us patterns or alerting us to anomalies. We imagine as we analyze the research data in *our own work settings*, in addition to computer vision, big data analytics and speech detection and interpretation systems that we could use to analyze the field data, we have means to realistically (not 'effortlessly', but not so difficult as to be not worth the effort needed) combine our time-and-place-specific data with time-and-space congruent 'big data' sets such as traffic conditions and social media postings.

So, what might meaningfully change in terms of first order consequences, i.e. how we conduct research? Given that we will have real time access to machine data, the main differences in how we sense and sense-make shift from happening exclusively after the fact/in the office, to a mix of in real time/in the field and in office. We imagined two significant outcomes for our field research practices.

We See and Hear in Field What We Couldn't Before – Observation and sensemaking by machines move closer together in time and space, with real-time, or close to real-time, pattern and anomaly detection happening in the field on our data capturing equipment rather than after the fact in our offices. We imagine this means our field equipment will alert us to movements, behaviors, activities, environmental conditions, and to the presence or arrangement of objects or to turns of phrase, language miscommunications and non-speech audio events that seem significant in the moment. In addition to simply noting and alerting us to patterns or anomalies, our equipment might also suggest actions to take such as a question to ask about an object as we unpack a car, a rearrangement of objects on our sorting sheet as we inventory the contents of a car, or an data analysis point and a suggested follow up question such as: *83.5% of the car contents belong to Fernando, but Ana Luisa is the primary user of this car: how did this come to be?* The frequency and intensity of these alerts and

suggestions will likely increase as a given interview progresses, and over the course of a project as the number of interviews we complete increases, as the equipment could be adding to its knowledge set over time depending on how learning and inferencing are architected.

Interviews Are More Exhausting for Ethnographers – Currently during fieldwork we are intensely and actively listening to participants and noticing everything around us, even as we formulate our next question and continually recalibrate the overall flow of a conversation with participants. This requires intense focus and is, frankly, exhausting. Indeed, on first exposure to field work during LEAM, our electrical engineer teammate marveled that the ethnographers could keep an enthusiastic conversation going with a participant for as long as we normally did. With sense and sensemaking smart machines in field with us, such work will require we expertly handle and incorporate an extra stream of information coming at us, and fluidly incorporate it into our orchestration of conversation and observation.

In our office locations, as we analyze data between field visits and at the conclusion of the data collection, we see one significant change.

We Better Distinguish and Flag Possible Differences in Causality of Data Patterns – We imagine being able to combine our GPS/telemetry or other machine-created data with other data sets so that we can understand our individual participant actions in the context of larger time- and space-specific events to highlight possible connections. In our original fieldwork, we almost missed an important story in Brazil when we mistook the GPS and phone data from a participant to indicate she had trouble parking (she made a phone call and she drove very slowly in a meandering pattern through a neighborhood). We brought biases to our interpretation of the data from similar data patterns we had seen elsewhere in Brazil and in Germany, and commented to the participant that her data from one Friday night seemed to show a hoh-hum evening of looking for a parking spot. The participant corrected us and explained that she had been alerted that night to an *arrestão* in the neighborhood (a group of criminals moving through an area and robbing everyone they encounter) by a phone call from a friend, she heard gunshots in the distance, and spent a frantic twenty minutes moving slowly out of the area hoping not to be robbed. Had we been able to run our data against social media postings or police reports from that neighborhood at that time, we could have been alerted to follow up with the participant in the second interview. We would, as well, get closer to the ‘ground truth’ our colleague imagined machines could detect.

Combined, these new field and office based sensing and sensemaking capabilities will shape how we create data and how we quickly or frequently we iterate our interviews or other research protocols. Given these outcomes, what might meaningfully change in terms of second order consequences, i.e. professional expectations for technology industry ethnographers? We see three additional outcomes beyond those outlined in the low level environments.

New Professional Training and Expectations For Handling Smart Equipment, Information, and Analysis Coming at the Ethnographer Real-Time in the Field – We expect such training opportunities will likely show up first in professional organizations, such as EPIC, where

professionals can update their already expert ethnography skills. Moreover, we don't imagine that the presence of sense and sensemaking equipment in field means companies can hire 'cheap interns' as the human work is reduced to description and guided by machines; rather such skilled human work will require more training and more experience to deftly combine human questions and analysis and suggestions from equipment.

New Expectations for Fluency in Collaborating with Data Scientists and Computer Scientists

– Underlying the examples of smart equipment we've given is an assumption that the software on the equipment is trained to analyze and to learn over time. We do not believe it likely that all ethnographers will be fluent in the broad range of intense work that goes into building AI solutions: collecting data, creating a model, tuning it, implementing it on equipment and maintaining the software and hardware over time. They will, however, be expected to know how to effectively partner with experts who can do such work and to work together to define the expected outputs from the smart equipment. Ethnographers will need some ability or sensitivity to 'think like a computer or data scientist' just as we currently have some sensitivity and ability to 'think like an engineer' (and not give them too many problems). We need to be able to conceptualize: how can we translate what we need into data requirements that our colleagues can use when building a solution to train a machine?

New teams, New Research Protocols and New Standards of Data Analysis – Tying together the first two secondary outcomes, is the larger employment context in which technology industry ethnographers will work. In a world with mid-level presence of advanced sensing and sense making capabilities in our research and work settings, ethnographers will need to work with computer and data scientists *before* fieldwork, as they collaboratively scope projects, research goals and protocols. Much as ethnographers made a shift in the 1990s and early 2000s to conducting 'digital ethnography', in the 2020s technology industry ethnographers will shift to truly working in, as well as studying, cyber-physical worlds.

Home Instrumentation and Sensing Study (HISS)

Our last case study will be our most extreme, as we imagine adding a high level of machine sensing and sensemaking into our field and office locations. For our Home Instrumentation and Sensing Study (HISS) completed three months ago, our goal was to revamp our existing Smart Home usage roadmap to explicitly include householders' domestic lives with a new roommate: artificial intelligence. Quite a bit has changed in three years: Intel now describes the Smart Home on our external corporate website as "perceptive, responsive and autonomous", a three-adjective shorthand for enabled by artificial intelligence. A vision video on the same site titled "Smart Homes Are like Us: real-time collecting | analyzing | diagnosing" illustrates usages requiring advanced sense and sensemaking capabilities working near real time in the home. (Intel 2018b). In our offices, we repeatedly find ourselves in conversations with computer scientists, engineers and other product team members who make assumptions about 'always on' sensing in homes, often through cameras (coupled with computer vision algorithms and other capabilities) and microphones (coupled with automatic speech recognition, natural language processing, acoustic event detection and other capabilities). As social scientists, we find these conversations alarming and fascinating

— in other words, urgently in need of data in-and-of-the-world to validate or disprove the assumptions being made about the necessity and desirability of always-on sensing in homes.

With the HISS research protocol we honed in on exactly how, where, when and why householders might (and might not) want their homes to be perceptive, responsive and autonomous although we kept our research protocol light-hearted and free of such jargon or any mention of machine learning, computer vision, artificial intelligence or even the term smart home. We will not describe the entire protocol here, but rather limit ourselves to the parts that could be most radically altered if a high level of advanced sense and sensemaking systems were already widely present in American homes.

In HISS, we used a smart phone ethnography tool to have 101 US householders create a corpus of scenarios for living in imaginary versions of their own homes, specifically ones instrumented with sensing and inferencing capabilities that could make novel experiences possible. Participants started with two assignments that allowed us to understand what they valued in their homes now. With videos and text answers, participants shared details around the three areas of their homes in which they spend the most time awake. They then detailed two changes they would make to their homes if they won a “Complete Home Makeover With an Unlimited Budget”, and explained why they chose each change and how each would affect their home lives. We then shifted participants to a series of five assignments that playfully teased out householders’ overall expectations for the experience of living in an AI-enabled home, where and when they might want to partner with AI to a specific end, how an AI enabled home should know what’s happening in and around it, and how they expect an AI home should function, including their expectations for the accuracy and consistency of the sensing and inferences made. How did we accomplish this tall order? We asked participants to imagine winning a “free upgrade” to their home makeover so that their homes could always-and-all-over be able to sense and make sense based on one input modality. First they imagined life in a “smell-o-matic” version of their homes, then, (alternatively, and in turn) “sight-o-matic”, “touch-o-matic”, “hear-o-matic” or “taste-o-matic” versions. In each assignment, householders shared how they use their human sense at home now (“I walk in the door and smell dinner cooking and I know we’ll be eating kimchi stew for dinner”), acted out in video what it would be like to live in a sensing home, and then in text gave us an additional three “if this than that” style scenarios in which their home senses objects, actors or activities, then infers and acts to change their home lives. We also asked participants for a scenario in which sensing in their homes would not be welcome⁹. Moving participants through the staged protocol, checking the quality and completeness of their answers, asking clarifying questions, or for re-dos was a full time job (weekends, evening, and holidays) for us. We felt, at times, overwhelmed by the data coming in and the need to keep participants moving.

As we only recently finished the data analysis, the concrete outcomes of HISS are still developing. We are partnering with product teams to apply the research insights to the definition of smart home AI prototypes and products. Moreover, HISS was planned as the first of a series of projects to create AI experience roadmaps in different usage contexts and we are moving forward with other research¹⁰.

HISS in a Sense and Sensemaking Saturated Environment

Let's now imagine HISS with a high presence of advanced sensing and sense making technologies in our research participants' homes and in our corporate offices. In this alternative reality, high presence in our field locations might mean that householders already:

- live in homes with sensors, devices, connectivity and computing solutions that makes it possible to:
 - detect odors and signature scents in the air (i.e. smell)
 - detect and decode vibrations that travel through the air (i.e. hear)
 - detect and decode meaning through tactile sensations like pressure, hot and cold, wet and dry, and vibration (i.e. touch)
 - create a two or three dimensional likeness of the environment (i.e. see)
 - detect chemicals and substances in liquids or solids (i.e. taste)
- live in homes that have connected infrastructure, appliances and devices that can respond to conditions and commands from humans or machines, and can make changes in the home (actuators and automation)
- are fluent in interacting with an in-home and on-phone interface that lets them direct, respond to, and experiment with new usages of this system. The popularity of smart speakers - recent research predicts almost 50% of US households will have a smart speaker by 2019 (Adobe 2018) – makes this at first far-fetched scenario not so unimaginable.

In our work locations, we imagine we have a means of securely accessing data and event logs from participants' homes during the course of the study, from sources including utilities, internet providers, and of course the sensors, devices and computing solutions within the homes.

So, what might meaningfully change in terms of first order consequences, i.e. how we conduct research? We imagine four significant outcomes to our practice.

Machines Co-Design Research with Ethnographers – We promised we would not venture to imagining sentient robots conducting fieldwork, and as we are re-envisioning a remote ethnography study there will be no need to do so¹⁰. Instead, inspired by Autodesk's much heralded 2017 Elbo chair created as a collaboration between a human designer and Dreamcatcher software, we imagine ethnographers similarly collaborating with AI software. The future ethnographer may input parameters for research goals, participant criteria, data types and data reliability, research cost and duration constraints, and have an AI system trained with details and successful/failures, impacts and outcomes of earlier projects produce multiple research plans for remote ethnography that can be further tweaked as the ethnographer provides more inputs. Can an AI system write more creative and engaging mobile ethnography questions than we did for HISS? Certainly we expect it could be faster, as we spent a *generous* amount of time designing our questions. Even if the system could provide variations on our existing questions, this could save us weeks of time.

Machines Co-Execute Research with Ethnographers – In a high presence of advanced sensing and sense making technology world, we imagine machines can simplify recruiting

and managing research participants in at least three ways. First, gone would be traditional screeners we use now; in would be our machines matching potential participants to our research criteria based on the data from their already instrumented homes that householders have consented to provide for commercial purposes. Our machines could contact these participants through their in-home systems or smart phones and ask if they would like to participate, and gather further information from them so that the human ethnographers can review and choose the most suitable participants, saving us weeks of work. Second, as participants are accepted into the study, our machines could negotiate access to relevant and limited home systems and secure informed consent from participants about how and when the data will be collected, stored, analyzed and deleted. Third, as the research goes to field, the ethnographers could rely on our machine partners to check video, audio and text responses for quality based on criteria we provide. If responses are unclear or inadequate, the system could generate follow up questions or instructions, and only alert the ethnographers when a resolution to inadequate data cannot be negotiated. Overall this could easily save us 100 hours of work over the course of a project the size and scope of HISS.

Ethnographers More Often Experiment and Evaluate Using Data to Create Change In the World – Recall that as social scientists we assume our job is to generate data to identify opportunities to make change in the world. In a high presence world this changes. We likely more often than not end up guiding participants to generate and use data to create change in their own worlds, moving us as ethnographers closer in research practice to our computer science and engineering colleagues. While as ethnographers in technology industry we have certainly engaged in evaluative research from concept testing to proof of concept prototype evaluations to alpha and beta testing of product releases, what we propose here is novel type of evaluative intervention. Rather than just imagining and play acting scenarios in HISS prompted by our research questions, future participants could detail what they would want to be possible and their home systems would attempt to deliver the experience. In part, our research data would then consist of the experiences participants proposed to their systems and the data inputs they ask the systems to consider; the choice their home systems make to meet experience and data criteria; the success, failure and reactions of participants to the new home capabilities; and the negotiating and updating participants would do with their homes systems to achieve more satisfactory results. We would still have participants reflect with us on the value of the usage; combined with the above data, we would indeed be much closer to the ground truth data inputs that our colleagues need to build better AI solutions.

Mediate Better Transparency Between Established and Emerging Experts in Cyber-Physical Worlds – Currently we think of computer and data scientists, engineers, and developers as experts on machine learning, computer vision, natural language and audio processing, and Internet of Things sensing and connectivity technologies. In the future, we will also consider people who work, move through, live, shop, recreate and more in environments with high presence of sense and sensemaking technologies to be experts on these systems. These people will be knowledgeable and skilled at interacting with, training, working around, and occasionally subverting such systems. As ethnographers are adept at handling data *about people* and will be adept with data *traces by people* we will play an important role in tech companies deploying such systems. Through our research and advocacy, we will make the

priorities, concerns, desires and actions of both types of experts more transparent to one another. Ethnographers will champion ethics, accountability, and data rights to guide strategies and policies their companies use in solution development. Given the ability of these systems to exercise real power in people's lives, having ethnographers who can facilitate understanding and empathy for the positions of all experts will be critical for adoption of solutions that do not create a total surveillance society. Returning to Paglen's (2016) critique of the deployment of computer vision systems as "an active, cunning, exercise of power, one ideally suited to molecular police and market operations—one designed to insert its tendrils into ever-smaller slices of everyday life" we will have an obligation to shape how that power is used. As ethnographers, our deliverables will need to speak to much broader and much more senior audiences at our companies; we will no longer be a nice to have, or anecdotal or fluffy; but a crucial part of mature AI solution development practices.

Given these outcomes, what additional changes to the professional expectations for technology industry ethnographers might follow? We see two additional outcomes beyond those outlines in the low and medium level environments.

Future Ethnographers are Trained Differently –Ethnographers must be fluent in querying how digital traces are created, and in assessing what they can and cannot understand from them. They must also develop new methods for engaging research participants skilled in living in ever more algorithmically generated and mediated settings and fieldsites that can rapidly be reconfigured based on input from lay and professional experts. These skills will become standard part of training at universities as well as professional development courses. Australia National University's 3A institute, with a charter to deliver "a new applied science to enable the safe, ethical and effective design, integration, management and regulation of cyber-physical systems" is one example of where we already see happening in education. (ANU 2018)

Ethnographers Become More Visible at and More Important for Accountable Technology Companies – Future ethnographers, as experts who sit comfortably at the intersection of the cyber and physical worlds, will share peer stature with the data scientists who have come to the fore in the new data-driven innovation narrative popular in technology companies. With complementary skills, they will be jointly responsible for evaluating and justifying the types of sense and sense making capabilities their companies enable, that in practice can quickly devolved into the surveillance and exercises of power Paglen notes. Ethnographer's stars will rise at *accountable technology companies*; companies that expect to create products explainable to those whose lives are shaped by them and who expect the actions and decision these systems make to them. Ethnographers working for unaccountable technology companies can expect to battle to shift their companies to accountability, or be reduced to apologists for controversial exercises of power using invisible technologies.

CONCLUSION

Through re-imagining three of our existing ethnographic research projects, we have attempted to draw reasonable expectations for how ethnographic practice in the technology

industry may change in the next 20-25 years. Our explorations have been unapologetically ethnography-centric. We did not choose to imagine a dystopian future in which the ethnographer's expertise and skills in engaging with, understanding, and thoughtfully giving voice to people outside tech will lose all currency with our employers. Rather, we provided specific and rather pedestrian examples of how we might integrate new technology capabilities into our existing practices without completely blowing up the concept of applied ethnography as a practice and an approach to sensemaking grounded in social science theory and methods. We realize that we have, at times, taken great liberties in our imagining; we have glossed over the complexities, for example, in building machine learning systems that could perform the types of anomaly and pattern recognition we have imagined. These complexities including sourcing suitable data for training the system, and the time and cost in developing such a solution, among many others. In this way, as well, we have been ethnography-centric, choosing to focus on how the impressive and complex work undertaken by our data scientist, computer scientist, engineering, software programming and other highly trained technology industry colleagues could serve *our* professional ends.

We have argued that the combination of technology capabilities that together are common glossed as AI, including machine learning, Internet of Things, computer vision and speech and audio processing technologies, do offer compelling new ways to model and understand the context in which ethnographic encounters take place. They will most likely become indispensable in allowing ethnographers to see and hear data we may not realize we have. They can help us interpret our specific data in the context of larger events or larger patterns. In both ways, they expand our capabilities as researchers and allow a space for us as ethnographers to engage in the new data-centric innovation narrative in ways that acknowledge our expertise in understanding people, and provide an avenue for us to influence how such technologies are developed over time.

In some fashion adding these capabilities to our research practice does get us closer to ground truth about some human behaviors; they do offer better representations of the things we set out to measure. But as we argue that ethnography has never been about testing hypotheses to ultimately generate laws, these new technology tools will never get us *on their own* to the type of truths the ethnographer values above all else; the meanings given to experiences by humans, meanings for which the ground beneath is always unstable. We look forward to the next twenty years, to see how our practices change (and what we got right and wrong) and to the excitement and creative frictions that future rumbles, jostles and liquefaction of our grounds for building evidence will produce.

Alexandra Zafiroglu is a principal engineer at Intel. As a cultural anthropologist, she immerses herself in others' daily routines and realities and formulates explanatory and predictive models for how people and organizations create and will respond to techno-economic transformations in a data-centric world. Contact her at: zafiroglu@gmail.com

Yen-ning Chang is a design researcher at Intel. With a background in Journalism and HCI, she is passionate about communicating people's stories. She brings concepts to life by creating empathy within multi-disciplinary teams to ensure the design work is underpinned by a strong human-centered foundation. Contact her at: yenning.c@gmail.com

NOTES

The research projects explored here were sponsored by Intel Corporation. The arguments presented in this paper do not represent the official position of Intel. We thank the range of research partners for these projects, including: Jennifer Healey, Tim Plowman, David Graumann, Georgios Theocharous, Philip Corriveau, Susan Faulkner, Laura Rumbel, Kathi Kitner, Heather Patterson and Faith McCreary.

1. https://en.wikipedia.org/wiki/List_of_technology_centers
2. Here Woods is paraphrasing Hilary Mason (Mason and Wiggins 2010)
3. Most writing on this topic is a mix of each of these arguments, rather than the playful ‘stages of grief’ dissection that we’ve done here. Examples include: Anderson et al 2009; Boyd and Crawford 2012; Burrell 2016; Elish & Boyd 2017; Madsbjerg 2017; Wang 2016.
4. We recognize a number of critiques of this model, including Dourish’s (2006) oft-cited critique of the ‘implications for design’ model, and Amirebrahim’s (2016) critique of the flattening of social science through user experience. While these critiques are important, they are not our focus here.
5. Rather than name drop and speak inside baseball, we tend to use key concepts backed up with specific examples. If stakeholders want to know more, we gladly share – but we never overtly start with theory. Elisabeth Shove’s definition of ‘convenience’ in *Comfort, Cleanliness + Convenience: the social organization of normality* (2003) was critical in explaining the history and future of the smart home to stakeholders in our first case study in this paper. Jenna Burrell’s 2016 paper on machine learning has been a useful piece for helping us explain the differences between human and machine cognition, and how we ask question ‘about AI’ to non-experts in our last case study in this paper.
6. Paraphrased from memory; presented at an invited talk at NWWiC Regional Conference 2013 (Zafiroglu & Healey 2013).
7. The work of Boyd & Crawford (2012), Kitchin (2014), Wang (2014) and Boyd & Crawford (2017) as well as the collective body of work published in *Big Data & Society* amply makes these points.
8. A perspective on the role of ethnographers’ responsibilities for designing for privacy in smart homes based on SHIFT was presented at EPIC 2016 (Zafiroglu, Patterson, McCreary 2016).
9. The entire protocol is longer and more detailed than is relevant in this case study. We are currently working on a conference paper that details HISS and findings.
10. The next planned project will be for AI in manufacturing, and will, naturally, be named MISS for ‘Manufacturing Instrumentation and Sensing Study’.
11. For those of you craving a good robot-ethnography story, Alicia Dudek’s delightful “Lou and Cee Cee prepare for fieldwork in the future: a world where robots conduct ethnography” (2016) will not disappoint!

REFERENCES CITED

- Adobe Digital Insights
2018 State of Voice Assistants First published September 10 2018
<https://www.cmo.com/features/articles/2018/9/7/adobe-2018-consumer-voice-survey.html#gs.CcIsBMW>
- Amirebrahim, Shaheen
2016 The Rise of the User and the Fall of People. *Ethnographic Praxis in Industry Conference Proceedings* 2016:71–103.

- Anderson, Ken, Dawn Nafus, Tye Rattenbury and Ryan Aipperspach
 2009 Numbers Have Qualities Too: Experiences with Ethno-Mining. New York: *Ethnographic Praxis in Industry Conference Proceedings* 2009:123–140.
- ANU
 2018 3AI Institute Home Page accessed September 20 2018 <https://3ainstitute.cecs.anu.edu.au/#about>
- Bogost, Ian
 2017 Why Zuckerberg and Musk are Fighting About the Robot Future *The Atlantic*, July 27 Accessed September 1, 2018. <https://www.theatlantic.com/technology/archive/2017/07/musk-vs-zuck/535077/>
- Bowley, Rachel
 2017 The Fastest-Growing Jobs in the U.S. Based on LinkedIn Data first published December 17, 2017 <https://blog.linkedin.com/2017/december/17/the-fastest-growing-jobs-in-the-u-s-based-on-linkedin-data>
- Boyd, Danah and Kate Crawford
 2012 Critical Questions for Big Data *Information, Communication, & Society* 15(5):662-679
- Burrell, Jenna
 2016 How the Machine “Thinks”: Understanding Opacity in Machine Learning Algorithms. *Big Data & Society* first published January 6, 2016 <https://doi.org/10.1177/2053951715622512>
- Davenport, Thomas H. and D.J. Patil
 2012 Data Scientist: The Sexiest Job of the 21st Century. *Harvard Business Review* 90(10):70-76, 128
- Dourish, Paul
 2006 Implications for Design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* New York: ACM 541-550.
- Dudek, Alicia
 2016 Lou and Cee Cee prepare for fieldwork in the future: a world where robots conduct ethnography *Ethnography Matters* first published June 17, 2016 <http://ethnographymatters.net/blog/2016/06/17/lou-and-cee-cee-prepare-for-fieldwork-in-the-future-a-world-where-robots-conduct-ethnography/>
- Elish, M.C. and danah Boyd
 2017 Situating Methods in the Magic of Big Data and AI. *Communication Monographs* 85(1):57-80
- Feltman, Rachel
 2014 MIT Researchers Can Listen to Your Conversation by Watching Your Potato Chip Bag. *Washington Post* August 4, 2014. https://www.washingtonpost.com/news/speaking-of-science/wp/2014/08/04/mit-researchers-can-listen-to-your-conversation-by-watching-your-potato-chip-bag/?utm_term=.ba5473c31113
- Geertz, Clifford
 1973 *The Interpretation of Culture: Selected Essays* New York: Basic Books
- Glassdoor
 2018 Fifty Best Jobs in America Accessed September 1, 2018. https://www.glassdoor.com/List/Best-Jobs-in-America-LST_KQ0,20.htm
- Grace, Katja, John Salvatier, Allan Dafoe, Baobao Zhang and Owain Evans

2018 When Will AI Exceed Human Performance? Evidence from AI Experts. *Journal of Artificial Intelligence Research* 62: p. 729-754 <https://arxiv.org/abs/1705.08807>

Intel Corporation

2018a Media Alert: Data-Centric Innovation Summit – Data Center Platform and Products Fueling Intel's Growth <https://newsroom.intel.com/news-releases/data-centric-innovation-summit-data-center-platform-products-fueling-intels-growth/> (August 7 2018)

Intel Corporation

2018b Smart Homes: Perceptive, Responsive, and Autonomous accessed September 1 2018 <https://www.intel.com/content/www/us/en/smart-home/overview.html>

Kitchin, Rob

2014 Big Data, New Epistemologies and Paradigm Shifts. *Big Data & Society* first published April 1, 2014 <https://doi.org/10.1177/2053951714528481>

Madsbjerg, Christian

2017 *Sensemaking: The Power of the Humanities in the Age of the Algorithm* New York: Hachette.

Mason, Hilary and Chris Wiggins

2010 A Taxonomy of Data Science *Dataists.com* First published September 25, 2010 <http://www.dataists.com/2010/09/a-taxonomy-of-data-science/>

McKinsey

2009 Hal Varian on How the Web Challenges Managers *McKinsey Quarterly* January 2009 <https://www.mckinsey.com/industries/high-tech/our-insights/hal-varian-on-how-the-web-challenges-managers>

Metcalf, Jacob and Kate Crawford

2016 Where are the Human Subjects in Big Data Research? The Emerging Ethics Divide *Big Data & Society* first published June 1, 2016 <https://doi.org/10.1177/2053951716650211>

O'Neil, Cathy

2016 *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* New York: Crown Books

Paglen, Trevor

2016 Invisible Images (your pictures are looking at you) *The New Inquiry* December 6, 2016 <https://thenewinquiry.com/invisible-images-your-pictures-are-looking-at-you/>

Reichenbach, Lisa and Magda Wesolkowska

2008 All that is Seen and Unseen: The Physical Environment as Informant *Ethnographic Praxis in Industry Conference Proceedings* 2008:160-174.

Schwab, Klaus

2017 *The Fourth Industrial Revolution* New York: Crown Publishing

Shove, Elizabeth

2003 *Comfort, Cleanliness and Convenience: The Social Organization of Normality* Oxford: Berg

Wang, Tricia

2016 Co-Designing with Machines: Moving Beyond the Human/Machine Binary first published June 13, 2016 <http://ethnographymatters.net/blog/2016/06/13/co-designing-with-machines-moving-beyond-the-humanmachine-binary/>

Woods, Dan

2011 LinkedIn's Daniel Tunkelag on "What is a Data Scientist" *Forbes* October 11, 2011
<https://www.forbes.com/sites/danwoods/2011/10/24/linkedin-daniel-tunkelang-on-what-is-a-data-scientist/#2ab65e3311cc>

Zafiroglu, Alexandra, Jennifer Healey and Tim Plowman

2012 Navigation to Multiple Local Transportation Futures: Cross-Interrogating Remembered and Recorded Drives *AutomotiveUI '12: Proceedings of the 4th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* p. 139-146

Zafiroglu, Alexandra and Jennifer Healey

2013 "From Ground Truth to Common Ground: Negotiating Between Problem and Solution Spaces"
invited talk at NWWiC Regional conference Portland, OR, 2013

Zafiroglu, Alexandra, Heather Patterson and Faith McCreary

2016 Living Comfortably In Glass houses 2016 *Ethnographic Praxis in Industry Conference Proceedings*, p. 540, ISSN 1559-8918, <https://www.epicpeople.org/living-comfortably-glass-houses/>.

Zhang, Chao and Juchi Yan, Changsheng Li, Xiaoguang Rui, Liang Liu, and Rongfang Bie

2016 On Estimating Air Pollution from Photos Using Convolutional Neural Networks *Proceedings of the ACM on Multimedia Conference* Pages 297-301 doi>10.1145/2964284.2967230