**The Inference Engine: A Grand Challenge to Address the Context Sensitivity Problem in
Social Science Research**

Daniel L. Goroff¹, Neil A. Lewis, Jr.², Anne M. Scheel³, Laura Scherer⁴, & Joshua A. Tucker⁵

Sloan Foundation¹

Cornell University²

Technische Universiteit Eindhoven³

University of Colorado Denver⁴

New York University⁵

RUNNING HEAD: ADDRESSING CONTEXT SENSITIVITY IN SOCIAL SCIENCE

Main Text Word Count: 3,610

Author Note:

The ideas in this paper were generated through conversations among the authors at the 2018 Technical Exchange on Complex Social Systems at the Center for Open Science; order of authorship is alphabetical.

Abstract

Social science has a ‘context sensitivity’ problem: the people that we study, and the situations they engage in, are so complex and variable that predicting how they will think, feel, and behave in a given situation is very challenging. Even when we are able to make such predictions, it is often unclear how accurate they will be if some feature of the studied subjects and/or situation changes. This limits the utility of our research for application and policy, as the ‘contextual factors’ that might change our conclusions are often unknown. It is time to address this context sensitivity problem in social science research. While do not yet know *how* to solve it, we believe social scientists can make great progress by working together to build an inference engine.

Keywords: Context sensitivity, Replication, Causal Inference, Social Science

**The Inference Engine: A Grand Challenge to Address the Context Sensitivity Problem in
Social Science Research**

The upshot is that, if we want to know what would happen if we did X, we have to actually do X. Because we don't have inference tickets, we never get to take the shortcut. We always have to wade through the empirical morass. Always.

Borsboom, 2013

In 2003 Tony Tether, Director of the US Defense Advanced Research Projects Agency (DARPA), issued a challenge to create an autonomous vehicle; DARPA offered a \$1 million prize for whoever built a self-driving car that could drive 142 miles through the Mojave Desert the fastest (Davies, 2017). At the time, that seemed impossible; no one knew how to do it. In fact, the first time they issued this challenge, it failed miserably – every car crashed, failed, or caught fire. A few years later however, it was done. It took teams of engineers a few years to figure out how to do it, and now self-driving cars are roaming the streets of major cities all over the globe.

Today we are proposing a new challenge, this time for social scientists. By the year 2030, social scientists should have an “Inference Engine” that can, for example, take the results of a study run with 500 college students in Chicago and be able to accurately predict (e.g., within a 5% margin of error) what the results would be if that same study were run on 500 college students in Moscow, 500 senior citizens in Tokyo, and 500 adolescents in Kampala. Like Siri, Alexa, or Google Home, one should be able to ask this engine a question or request a prediction and get an accurate answer: “Ok Sanjay, look at the method and results of the Yeager and

colleagues (2016) lay-theory academic intervention and tell me what would happen if we ran the same thing in Shanghai”. Even better would be an engine that could respond “Ok Simine, if you want to produce an effect of the same size in Shanghai as Yeager and colleagues (2016) found in the United States, you will need to change the procedure so that the intervention is administered by a teacher face-to-face rather than through the internet” (or whatever the actual important contextual variable is).

To accomplish this, the engine would need to be able to pull information from all relevant previous studies on lay theory interventions, all information known about students who participated in the studies in the United States and how they might be similar or different from students in Shanghai, all relevant information about similarities and differences between Shanghai and the cities from the original studies, and estimates for what variables would change the size and magnitude of the effect in the new context. For a host of reasons summarized by Soderberg (2018), this is, at present, an impossible task. But like the driverless car example we opened with, we believe it is a current impossibility that, with coordinated effort among social scientists, can become a future reality.

What Would Be the Virtue of an Inference Engine?

If you have read this far you may be wondering why social scientists should devote time and energy to this “blue sky idea.” For social science research to be *useful* – to do more than fill journal pages – it is critical for us to learn which ‘contextual factors’ matter for our effects (Soderberg, 2018). People who are trying to use our research (including other researchers trying to build on it) need to know how different environments influence a phenomenon, the types of independent and dependent variables that have been examined and what they imply for both theoretical and practical knowledge, the attitudes, goals, or identity factors of participants that

have been examined and what those might mean, the time course of the phenomenon, which agents of change (e.g., doctors, teachers, government officials) might matter for how the study is interpreted by participants, and potentially much more (Earl & Lewis, in press).

To make this more concrete, consider the following examples: if a school district superintendent wants to implement the latest and greatest psychological intervention, they need to know whether their implementation plans need to be adjusted depending on class sizes within schools, the racial and/or socioeconomic composition of the classrooms, etc. If a clinician in Charlottesville, Virginia learns of a new behavioral therapy for racial minority patients that was developed in New York City, they may need to consider how differences in the racial climate of the cities may affect the quality of the doctor-patient interactions that may be essential for the therapy's success (e.g., Burgess, Warren, Phelan, Dovidio, & Van Ryn, 2010). In the military domain, a general might want to know how she may take information about how Germans reacted to the US opening a base in Heidelberg and use it to predict how Iraqis might react to opening a permanent base in Baghdad. These are not only practical questions – being able to answer them would also force us to generate better, more precise social scientific theories (Rothman, 2004).

Building the knowledge and infrastructure necessary to achieve the goal of creating a functional inference engine could lead us towards conducting research that is more efficient with respect to the generation, expansion, specification, and falsification of social scientific theories (see also Simons, Shoda, & Lindsay, 2017). That kind of research would ultimately be more useful for addressing social problems (Lewin, 1946). Finally, it would help social scientists (and those who fund us) with the prioritization of research: the engine could let us know in which areas our knowledge is stronger vs. weaker and hence need more or fewer additional studies,

how much additional resources are necessary to translate a finding to a new context/application, and help us predict the probability of success in a new research venture. Since time, resources, and attention are limited, this would help us to ensure that we are optimizing our use of limited resources as we advance our knowledge, ensuring that successive studies truly build upon our knowledge base in the places where more information is needed rather than, to borrow from Forscher (1963), adding more chaos in the brickyard that is social science research.

What is Currently Stopping Us?

A valid question to ask is why the inference engine we envision does not already exist. In the following we provide a non-exhaustive list of factors that may have hindered such endeavors in the past, and that may pose challenges for our goal to predict generalizability today.

Reproducibility is a precondition for generalizability. First, it is important to highlight the relationship between generalizability and replicability. Failures to replicate a result in a new context can occur because of the changed context or because the original result was a false positive (or an overestimation) and would fail to replicate even in the original context (Maxwell, Lau, & Howard, 2015). In other words, basic replicability – meaning the robustness of an effect in one context – is a precondition for generalizability. This is particularly important to consider because large-scale direct replication projects of the last years have shown that a substantial number of findings (between one and two thirds) in several social science disciplines cannot be replicated by independent authors (Open Science Collaboration, 2015; Camerer et al., 2016; Camerer et al., 2018). If up to two thirds of the published literature cannot be replicated using methods similar to those used by the original authors, relying on this literature to predict generalizability to new contexts will be an inefficient endeavor. Put differently, a lack of

replicability of original research adds a large amount of noise to the input of our inference engine, which will severely constrain the maximal level of output accuracy that can be reached.

Lack of dedicated research on generalizability. Additionally, there have been few studies that develop models for the generalizability of empirical findings to new contexts. Pearl and Bareinboim (2014) introduced a formal representation using graph-based procedures for deciding whether causal effects in a target population can be inferred from existing experimental findings. Their paper focuses on developing assumptions that must hold in order for the results to be “transportable” to another context. Suppose, for example, that a randomized trial occurs in Los Angeles that estimates the causal effect of exposure to a training program on employment status for different age groups in the LA labor force. If we wish to generalize the results to the population of New York City, suppose we know that the average age of the NYC labor force is significantly higher than in LA. Assuming that age-specific effects are invariant across cities, Pearl and Bareinboim (2014) develop a “transport formula” that combines the experimental estimate obtained in LA with observations about NYC’s population to generate an experimental estimate for NYC. In order to be useful, the methods developed in this paper require that the researcher has the necessary background knowledge to determine the dimensions along which two populations differ from each other.

What Tools Do We Currently Have At Our Disposal?

While the scope of this project – to be able to draw on all existing social scientific research studies and use it to make precise predictions about future studies – is far beyond anything currently in progress, there are some existing tools and ongoing research efforts that can help to facilitate our current goal.

Extracting relevant information from existing research

In recent years, meta-scientists have made a great deal of progress in training machines to extract information from previously published work. For instance, Nuijten and colleagues (2016) developed an algorithm – *statcheck* – that is able to automatically extract statistics from published research articles and recompute their *p*-values to check if there are serious errors in the published findings. Insights from that algorithm and others like it may be useful for extracting relevant information from existing research and making it more useful for future predictions. Efforts like the automation of statistical checking combined with more unified frameworks for evaluating the cumulative state of our knowledge (e.g., LeBel, McCarthy, Earp, Elson, & Vanpaemel, 2018) and standardizing practices for meta-scientific reviews (e.g., Arkesy & O'Malley, 2005) will more cleanly facilitate the development of the inference engine. We are encouraged by the fact that there is already work being done that is related to these tasks, in particular by people working on new approaches to conducting meta analyses (e.g., Bergmann et al., 2018; Bosco, Steel, Oswald, Uggerslev & Field, 2015).

The meta-scientific efforts described above will help to train the inference engine to be able to read scientific publications and extract relevant information from studies. At the most general level, the inference engine would need to know the *general area of inquiry*, the *specific research question*, the *research design*, and the *results* of the study. However, given that our focus is on understanding the *context* of different studies, we would also want to extract information that is going to be most theoretically relevant for understanding why a study carried out in one context might differ from a study in another context. While we expect that the contributions from individual research teams would be in part to identify exactly what these factors are, we could expect the list would likely include features such as the *age*, *sex*, and *ethnicity* of study participants, the *country* and *region* in which the study was located,

characteristics of the *time period* (e.g. seasonality, year, day vs. night, etc.) in which the study occurred, and *the mode of contact* (e.g., an intervention in a school, questions on a face-to-face survey, analysis of aggregate data) between the researchers and subjects.

What Else Would We Need?

While there are probably technological components of the inference engine that are yet to be discovered, we see at the current time three major sub-challenges that need to be addressed for this effort to be successful: (1) Building a pipeline to catalogue and inventory results of all existing social scientific research studies that is both retrospective (to draw on already existing studies) and can operate in real time (to draw on and integrate new studies as they are published); (2) Creating a global panel of all potential future subjects for research based on existing digital trace data currently produced by commercial, research, and government entities; and (3) Developing new machine learning models trained on existing research studies to ultimately predict the effects of interventions in context in which studies have not yet been conducted. Fortunately, research could begin on all three of these goals simultaneously, with the possibility of different research teams tackling different goals.

The *retrospective* portion of this task would involve designing machines that can extract information from existing archives of studies. We suspect that a great deal of this work would involve methods to work with PDF files of journal articles, but it may also involve ways of interacting with existing research archives such as The Open Science Framework (OSF), the Inter-university Consortium for Political and Social Research, the Harvard Dataverse, or the Roper Center for Public Opinion Research. The *real time* aspect of this research would involve designing (and advocating for the implementation of) new protocols for the use of meta data when articles and books are published in the future so that their findings can be seamlessly

incorporated into the inference machine's data infrastructure (see also Gorgolewski et al., 2017). One question with which researchers in this area may want to wrestle will be whether the inference engine should rely on the findings from existing studies to inform its machine learning models (given the issues with the existing published literature; Ioannidis, 2005), or whether we may ultimately want to move to a system where the inference engine draws upon the underlying data from these studies in performing its analyses. That is, would the machine focus on conclusions in text and tables from researchers, or would the machine actually go to sites such as the Harvard Dataverse or OSF and download raw data sets? The latter option might hold the promise of more valuable analysis, but would raise significantly more difficult technical challenges, to say nothing of legal and ethical challenges.

A Global Panel of Research Subjects

Here, we proceed from the assumption that the more information the inference machine has about the actual participants in a prior study – as opposed to just the characteristics of the study highlighted in the previous subsection – the better the machine will do at extrapolating the likely results of the study if it were carried out in another context. Researchers that are carrying out planned interventions always have the option of asking their subjects for information as part of the research protocol, but the number and type of questions researchers can ask and the answers they can obtain are always limited, if only by time, to say nothing of issues related to recall, social desirability bias, or the appropriateness of asking particular questions (Schwarz & Oyserman, 2001).

Today, however, we live in a world where increasingly large amounts of individual-level data are being digitally recorded at an ever-increasing pace. Some of this digital information is in the public domain – such as a Twitter profile, a comment on YouTube, or a post on an online

message board such Reddit – while other information is decidedly private, such as health or education records. Some of this information is boolean with records that absolutely speak directly to the question at hand (e.g., is a respondent registered to vote as a Republican or not?), while some information can be inferred (e.g., understanding potential medical concerns from Google Search data or political preferences from following patterns on Facebook). Collectively, however, the amount of digitized information about most individuals is vast.

Imagine, then, a world where every potential subject in a study had the ability to *voluntarily* allow researcher to incorporate the *sum total* of all digitized information about her or him as part of the research effort. Getting to this point would undoubtedly present significant *technical, social, legal, and ethical* challenges. *Technically*, we would need a method of marrying various sources of digitized records in all sorts of different formats to create a single unique record for each individually. *Socially*, this would involve many commercial, for profit, enterprises to agree to voluntarily share their most valuable commodity – their data – in the interest of science. Individuals, too, would need to consent to researchers accessing this data. The Sloan Foundation has been supporting a number of Administrative Research Data Facilities (ADRFs), which act as intermediaries of data between organizations and researchers. One recent example is an initiative to obtain Facebook data for research use. In the spring of 2018, the Social Science Research Council (SSRC), in collaboration with a consortium of research funders, worked with Facebook to secure a data use agreement that facilitates responsible access to social media data by academic researchers. Another initiative in the works is a plan by David Lazer and Joshua Tucker to develop an ADRF that specializes in social media data. Their goal is to work with companies to obtain data from multiple social media sources and grant access to researchers for responsible and reproducible research. *Legally*, huge barriers would need to be overcome to

allow companies that collect individual level digital data to share it collectively, to say nothing of government entities. As difficult as all of these challenges are, perhaps the most difficult would be overcoming the *ethical* hurdles of building such a dataset and ensuring that its contents were not used for nefarious purposes against the interests of the individuals whose data are included in it (not to mention adapting systems – like institutional review boards or legal systems – to adjudicate whether use is in fact nefarious). The Sloan Foundation has been interested in this issue and making grants in this direction for a long time (Goroff, 2015). One of the most recent grants has been awarded to Cynthia Dwork, who is working on developing techniques for allowing researchers to extract statistics from large datasets while at the same time protecting the privacy of individuals' data. Moreover, it has also recently supported a special semester on the foundations and applications of data privacy research, hosted at the Simons Institute for the Theory of Computing, which will be held in Spring 2019. This special semester brings together social scientists and computer scientists for an entire semester to work on research problems related to data privacy.

To be clear, this sub-challenge also contains great scope for research on how to go from digital trace data (e.g., an individual's posts on Twitter) to measures that will actually be useful for the inference engine (e.g., an individual's level of education).

New Machine Learning Models

Should the first two sub-challenges be overcome, a great deal of research would still need to be done to build machines that are actually capable of taking the inputs to the model – every scientific study ever conducted, and all we know about every individual who participated in a previous study based on their digital trace data – and producing the desired outputs: the likely result of that study conducted in a different context; the context in which studies ought to be

conducted next so as best to advance scientific knowledge, etc. This paper is hardly the place to get into details in this regard, but we envision machines that are able to simultaneously learn about studies on the same topic, studies in the same context, studies in the same population of participants, etc. to make predictions. So to go back to our original example of an experiment run on 500 college students in Chicago and a desire to understand what results a similar experiment on 500 elderly citizens in Tokyo might produce (say, for example, on the effect of a training course on digital literacy), we might expect the machine to learn how well all interventions regarding digital literacy have replicated generally, how they have varied when run on younger and older population, but also how experiments in Chicago and Tokyo on all subjects have tended to be replicated. Further, as we accumulated more data on individual study participants, one might expect the machine to learn if there were peculiar characteristics of the subjects in the Chicago study that suggest their reactions might not be typical of the general population, or even to produce predictions about how the study in Tokyo might turn out if different types of participant were recruited to take part in the study. Doing this requires a literature that is far more cumulative and a research culture that is more cooperative than currently exists. For instance, it may require that many people work on the same problem, rather than having each research team working independently on their own problems in different parts of the proverbial brickyard (Forscher, 1963).

Let's Get Started

While the idea of a social science inference engine seems quite fantastical at this moment in time – although perhaps no more than a car that drove itself might have seemed 15 years ago – we believe that breaking the challenge down into these three specific subcomponents provides a viable roadmap for how initial stages of research might begin. Moreover, even if the inference

engine was to remain in the realm of fantasy for an extended period of time, we believe that research in each of these core areas – how to extract information from existing social science research, how to use digital trace data to improve the quality of our analyses of future experiments, and how to train machines to take account of the different contexts in which different studies have been conducted in order to speak to the generalizability of that research – will all be extremely valuable in their own right.

References

- Bergmann, C., Tsuji, S., Piccinini, P. E., Lewis, M. L., Braginsky, M., Frank, M. C., & Cristia, A. (2018) Promoting replicability in developmental research through meta-analyses: Insights from language acquisition research. *Child Development*. Advance online publication. doi: 10.1111/cdev.13079
- Bosco, F. A., Steel, P., Oswald, F. L., Uggerslev, K. L., & Field, J. G. (2015). Cloud-based meta-analysis to bridge science and practice: Welcome to metaBUS. *Personnel Assessment and Decisions, 1*, 3-17.
- Burgess, D. J., Warren, J., Phelan, S., Dovidio, J., & Van Ryn, M. Stereotype threat and health disparities: what medical educators and future physicians need to know. *Journal of General Internal Medicine, 25*, 169-177.
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., ... & Heikensten, E. (2016). Evaluating replicability of laboratory experiments in economics. *Science, 351*(6280), 1433-1436.
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T. H., Huber, J., Johannesson, M., ... & Altmejd, A. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour, 2*(9), 637.
- Davies, A. (2017, November). *Inside The Races That Jump-Started The Self-Driving Car*. WIRED. Retrieved from: <https://www.wired.com/story/darpa-grand-urban-challenge-self-driving-car/>.
- Earl, A., & Lewis, N. A., Jr. (in press). Health in Context: New Perspectives on Healthy Thinking and Healthy Living. *Journal of Experimental Social Psychology*.
- Forscher, B. K. (1963). Chaos in the Brickyard. *Science, 142*(3590), 339.
- Gorgolewski, K. J., Alfaro-Almagro, F., Auer, T., Bellec, P., Capotă, M., Chakravarty, M. M., ... & Eklund, A. (2017). BIDS apps: Improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods. *PLoS computational biology, 13*(3), e1005209.
- Goroff, D. (2015). Balancing privacy versus accuracy in research protocols. *Science, 347*(6221), 479-480.
- Ioannidis, J. P. A. (2005). Why Most Published Research Findings Are False. *PLOS Medicine, 2*(8), e124.

- LeBel, E., McCarthy, R., Earp, B. D., Elson, M., & Vanpaemel, W. (2018). A Unified Framework to Quantify the Credibility of Scientific Findings. *Advances in Methods and Practices in Psychological Science*, 1(3), 389-402.
- Lewin, K. (1946). Action Research and Minority Problems. *Journal of Social Issues*, 2(2), 34-46.
- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *American Psychologist*, 70(6), 487-498.
- Nuijten, M. B., Hartgerink, C. H. J., van Assen, M. A. L. M., Epskamp, S., & Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985-2013). *Behavioral Research*, 48, 1205-1226.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716.
- Pearl, J. and E. Bareinboim. (2014). External Validity: From Do-Calculus to Transportability Across Populations. *Statistical Science*, 29(4): 579-595.
- Rothman, A. J. (2004). “Is there nothing more practical than a good theory?”: Why innovations and advances in health behavior change will arise if interventions are used to test and refine theory. *International Journal of Behavioral Nutrition and Physical Activity*, 1(1), 11.
- Simons, D. J., Shoda, Y., & Lindsay, D. S. (2017). Constraints on generality (COG): A proposed addition to all empirical papers. *Perspectives on Psychological Science*, 12(6), 1123-1128.
- Soderberg, C. K. (2018, July 31). DARPA TECSS Literature Review.
<https://doi.org/10.31219/osf.io/thbgd>
- Schwarz, N., & Oyserman, D. (2001). Asking questions about behavior: Cognition, communication, and questionnaire construction. *American Journal of Evaluation*, 22(2), 127-160.
- Yeager, D. S., Walton, G. M., Brady, S. T., Akcinar, E. N., Paunesku, D., Keane, L., ... & Gomez, E. M. (2016). Teaching a lay theory before college narrows achievement gaps at scale. *Proceedings of the National Academy of Sciences*, 113(24), E3341-E3348.