Morals and Methodology

Konstantin Genin

Group Leader: “Epistemology and Ethics of ML”

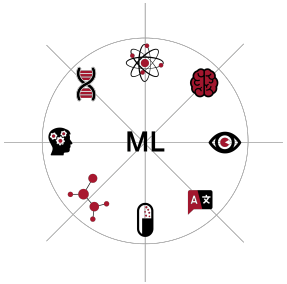

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Konstantin Genin

Understanding the (topological) complexity of scientific problems. Interactions between morals and methodology.

Vlasta Sikimic

Data-driven approaches to optimizations of scientific reasoning. Fairness and privacy in AI-assisted education.

Sander Beckers

Formalizing notions of actual causation, harm, responsibility, and explanation.

Sebastian Zezulka

Fairness in algorithmic prediction of long-term unemployment in the German context.

Morals and Methodology

Methodological questions depend on ethical questions.

- Unethical experiments ought not to be performed;
- Unjust algorithms ought not to be implemented.

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

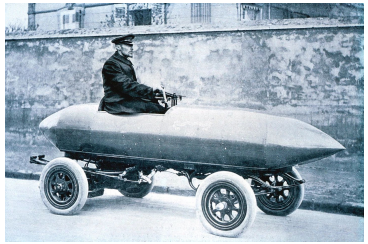

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

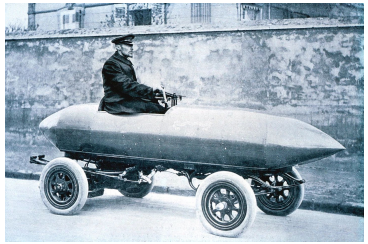

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

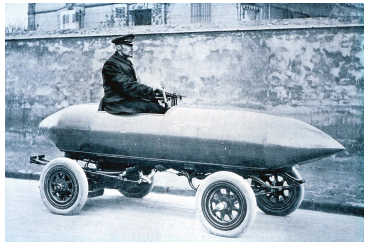

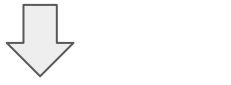

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

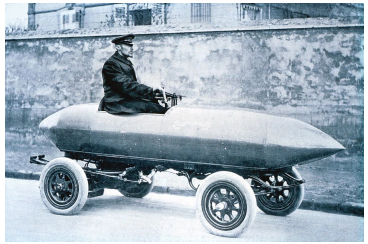

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

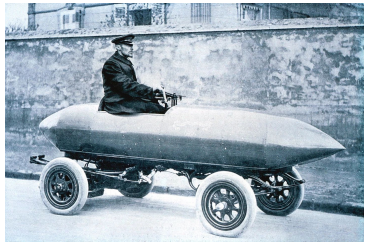

Morals and Methodology

As researchers advance methods, ethicists race to install new guardrails.

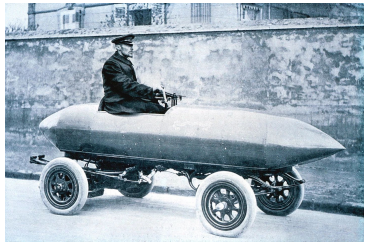

?

Morals and Methodology

But ethical questions also depend on methodological ones.

- RCTs in medicine and social science raise a variety of ethical issues, but are justified by their epistemic advantages, esp. for causal inference.

- As techniques for causal inference from observational data progress, there is pressure to reconsider the ethics of RCTs.

Morals and Methodology

Methodological progress **should** lead to moral progress.

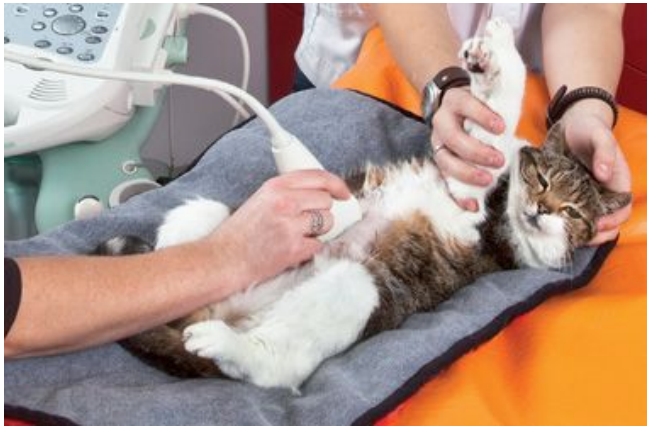

Morals and Methodology

The goal of the Ethics and Epistemology research group is to approach the issue from both sides:

- How exactly do methodological constraints express themselves as ethical constraints?

- How exactly should ethical constraints be expressed in methodology?

The Randomized, Controlled Trial (RCT)

“The RCT is the introduction of scientific method into the process of comparing treatments”

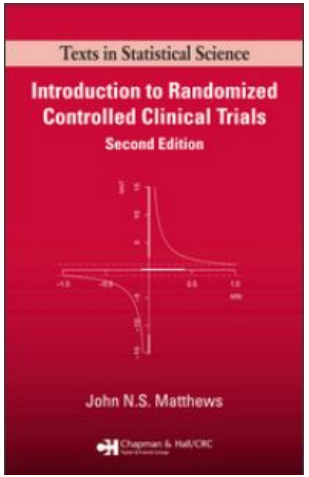

The Randomized, Controlled Trial (RCT)

Attempts to discover the relative effectiveness of a new intervention over standard treatment or placebo. Patients are assigned to the different “arms” of the trial by a randomization device.

- Widely considered the “gold standard” research design;
- Typically necessary for FDA, EMA approval;
- Raises a number of tricky ethical issues.

A call for RCTs for ML models in Clinical Settings

Perspective

Evaluating artificial intelligence in medicine: phases of clinical research

Yoonyoung Park[1], Gretchen Purcell Jackson[2,3], Morgan A. Foreman[1], Daniel Gruen[1], Jianying Hu[4] and Amar K. Das[1]

[1]Center for Computational Health, IBM Research Cambridge, Cambridge, Massachusetts, USA, [2]Center for AI, Research, and Evaluation, IBM Watson Health, Cambridge, MA, USA, [3]Department of Pediatric Surgery, Vanderbilt University Medical Center, Nashville, Tennessee, USA and [4]Center for Computational Health, IBM T.J. Watson Research Center, Yorktown Heights, New York, USA

Corresponding Author: Yoonyoung Park, IBM Research, 75 Binney Street, Cambridge, MA 02142, USA (yoonyoung.park@ibm.com)

Received 11 September 2019; Revised 15 May 2020; Editorial Decision 25 June 2020; Accepted 1 July 2020

ABSTRACT

Increased scrutiny of artificial intelligence (AI) applications in healthcare highlights the need for real-world evaluations for effectiveness and unintended consequences. The complexity of healthcare, compounded by the user- and context-dependent nature of AI applications, calls for a multifaceted approach beyond traditional in silico evaluation of AI. We propose an interdisciplinary, phased research framework for evaluation of AI imple-

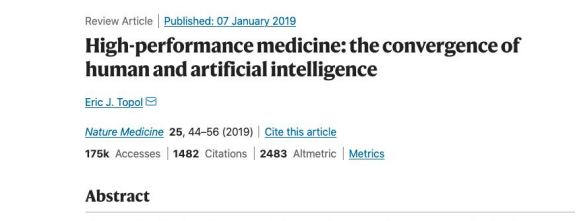

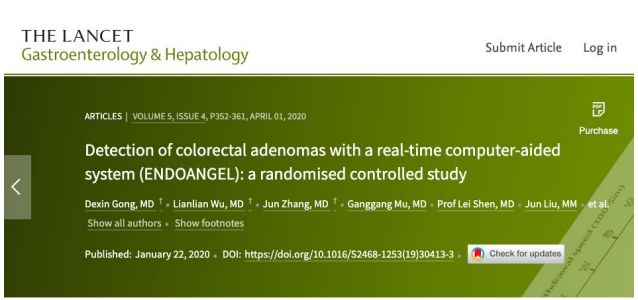

A Proliferation of ML methods for Causal Discovery

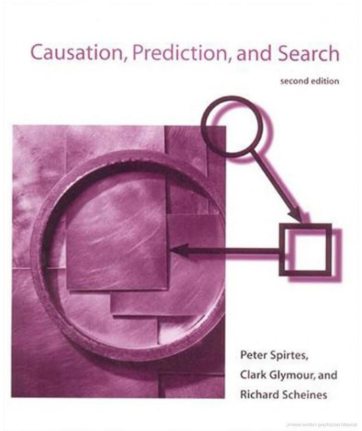

1993/2001

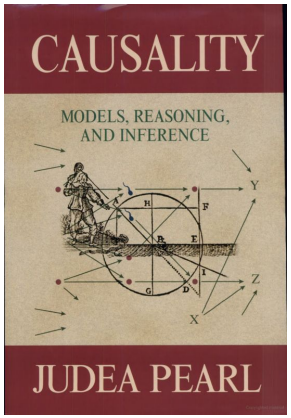

2000

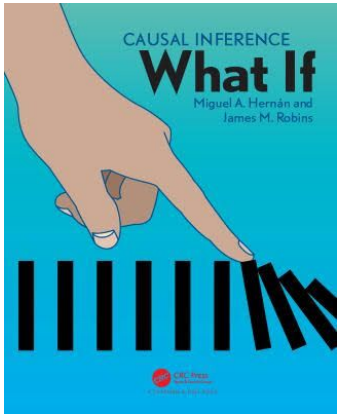

2010

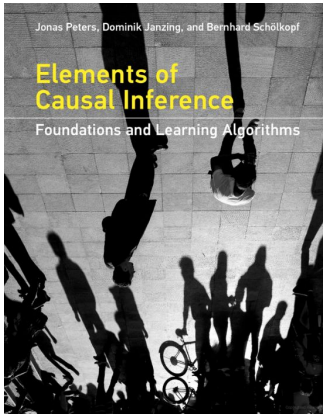

2017

A Proliferation of ML methods for Causal Discovery

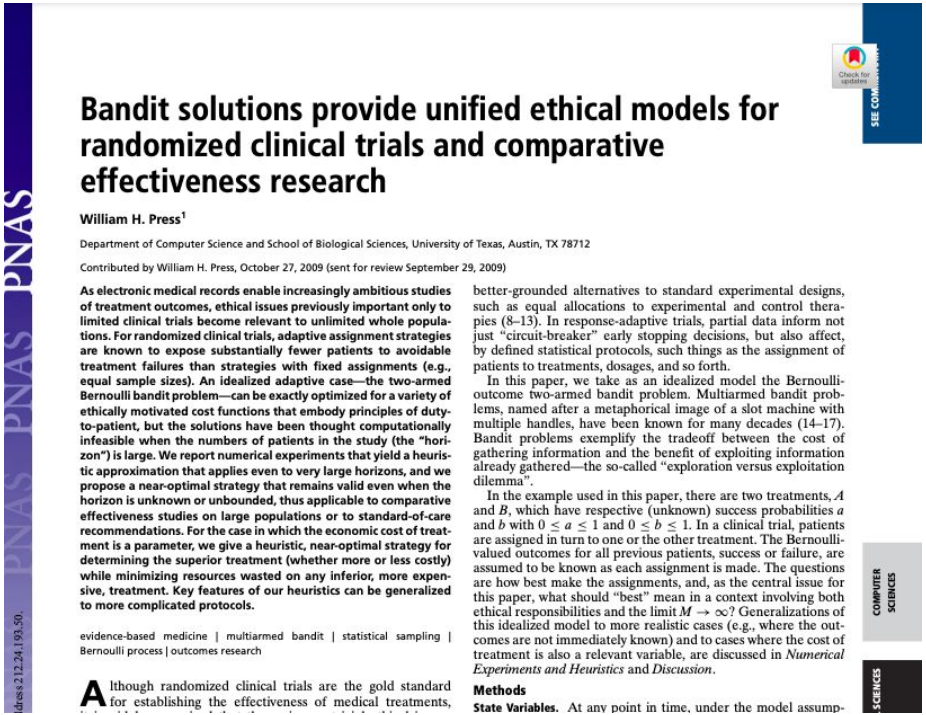

PNAS PNAS PNAS PNAS

SEE COMMENTARY

Bandit solutions provide unified ethical models for randomized clinical trials and comparative effectiveness research

William H. Press[1]

Department of Computer Science and School of Biological Sciences, University of Texas, Austin, TX 78712

Contributed by William H. Press, October 27, 2009 (sent for review September 29, 2009)

As electronic medical records enable increasingly ambitious studies of treatment outcomes, ethical issues previously important only to limited clinical trials become relevant to unlimited whole populations. For randomized clinical trials, adaptive assignment strategies are known to expose substantially fewer patients to avoidable treatment failures than strategies with fixed assignments (e.g., equal sample sizes). An idealized adaptive case—the two-armed Bernoulli bandit problem—can be exactly optimized for a variety of ethically motivated cost functions that embody principles of duty-to-patient, but the solutions have been thought computationally infeasible when the numbers of patients in the study (the "horizon") is large. We report numerical experiments that yield a heuristic approximation that applies even to very large horizons, and we propose a near-optimal strategy that remains valid even when the horizon is unknown or unbounded, thus applicable to comparative effectiveness studies on large populations or to standard-of-care recommendations. For the case in which the economic cost of treatment is a parameter, we give a heuristic, near-optimal strategy for determining the superior treatment (whether more or less costly) while minimizing resources wasted on any inferior, more expensive, treatment. Key features of our heuristics can be generalized to more complicated protocols.

evidence-based medicine | multiarmed bandit | statistical sampling | Bernoulli process | outcomes research

Although randomized clinical trials are the gold standard for establishing the effectiveness of medical treatments, better-grounded alternatives to standard experimental designs, such as equal allocations to experimental and control therapies (8–13). In response-adaptive trials, partial data inform not just "circuit-breaker" early stopping decisions, but also affect, by defined statistical protocols, such things as the assignment of patients to treatments, dosages, and so forth.

In this paper, we take as an idealized model the Bernoulli-outcome two-armed bandit problem. Multiarmed bandit problems, named after a metaphorical image of a slot machine with multiple handles, have been known for many decades (14–17). Bandit problems exemplify the tradeoff between the cost of gathering information and the benefit of exploiting information already gathered—the so-called "exploration versus exploitation dilemma".

In the example used in this paper, there are two treatments, A and B, which have respective (unknown) success probabilities a and b with $0 \le a \le 1$ and $0 \le b \le 1$. In a clinical trial, patients are assigned in turn to one or the other treatment. The Bernoulli-valued outcomes for all previous patients, success or failure, are assumed to be known as each assignment is made. The questions are how best make the assignments, and, as the central issue for this paper, what should "best" mean in a context involving both ethical responsibilities and the limit $M \to \infty$? Generalizations of this idealized model to more realistic cases (e.g., where the outcomes are not immediately known) and to cases where the cost of treatment is also a relevant variable, are discussed in *Numerical Experiments and Heuristics* and *Discussion*.

Methods

State Variables. At any point in time, under the model assump-

COMPUTER SCIENCES

SCIENCES

Check for updates

Article

Rethinking the Gold Standard With Multi-armed Bandits: Machine Learning Allocation Algorithms for Experiments

Organizational Research Methods
2021, Vol. 24(1) 78-103

DOI: 10.1177/1094428119854153
journals.sagepub.com/home/orm
SAGE

Chris Kaibel[1] and Torsten Biemann[1]

Abstract

In experiments, researchers commonly allocate subjects randomly and equally to the different treatment conditions before the experiment starts. While this approach is intuitive, it means that new information gathered during the experiment is not utilized until after the experiment has ended. Based on methodological approaches from other scientific disciplines such as computer science and medicine, we suggest machine learning algorithms for subject allocation in experiments. Specifically, we discuss a Bayesian multi-armed bandit algorithm for randomized controlled trials and use Monte Carlo simulations to compare its efficiency with randomized controlled trials that have a fixed and balanced subject allocation. Our findings indicate that a randomized allocation based on Bayesian multi-armed bandits is more efficient and ethical in most settings. We develop recommendations for researchers and discuss the limitations of our approach.

Keywords

experiments, randomized controlled trial, multi-armed bandit, exploration versus exploitation, machine learning, ethics in research

The Trouble with Randomization

Randomization comes into prima facie conflict with **therapeutic obligation:**

“A physician should not recommend for a patient therapy such that, given present medical knowledge, the hypothesis that the particular therapy is inferior to some other therapy is more probable than the opposite hypothesis” (Marquis, 1983).

The Trouble with Randomization

Randomization comes into prima facie conflict with **individualized treatment:**

“Although a patient who has been enrolled as a research subject in a RCT may benefit from the therapeutic effects of the treatment being tested, the fact that the treatment cannot be entirely tailored to that patient’s special needs seems to violate the physician’s obligation of unqualified fidelity to his patient’s health” (Schafer, 1983).

The Tragic View of Clinical Research

The discussion around clinical equipoise presupposes

- There is some valuable **epistemic good** secured by randomization;
- **Any** trial methodology which secures this good must **inevitably** come into conflict with the requirements of individual treatment**.**

The Tragic View of Clinical Research

The job of clinical ethics is to reconcile clinicians to this tragic situation:

“These clinical instincts, while understandable and laudable, have the potential to obscure the true nature of clinical research, as investigators and participants alike try to convince themselves that clinical research involves nothing more than the provision of clinical care. One way to try to address this collective and often willful confusion would be to identify a justification for exposing research participants to net risks for the benefit of others.” (Wendler, 2021).

The Tragic View of Clinical Research

The job of clinical ethics is to reconcile clinicians to this tragic situation:

“These clinical instincts, while understandable and laudable, have the potential to obscure the true nature of clinical research, as investigators and participants alike try to convince themselves that clinical research involves nothing more than the provision of clinical care. One way to try to address this collective and often willful confusion would be to identify a justification for exposing research participants to net risks for the benefit of others.” (Wendler, 2021).

David S. Wendler, M.A., Ph.D.

Senior Investigator

Department of Bioethics

NIH Clinical Center

VIEW SITE

Between Morals and Methodology

But is the tragic view **right?**

- What is the valuable **epistemic good** secured by randomization?

- Is there really **no** methodology that reconciles this good with the ethical requirements of individualized treatment?

Thank You!