

Assessing putative bias in prediction of anti-microbial resistance from real-world genotyping data under explicit causal assumptions

Mattia Prosperi

m.prosperi@ufl.edu

Department of Epidemiology, University of Florida
Gainesville, Florida, USA

Christina Boucher

cboucher@cise.ufl.edu

Department of Computer Science and Information and
Engineering, University of Florida
Gainesville, Florida, USA

Simone Marini

simone.marini@ufl.edu

Department of Epidemiology, University of Florida
Gainesville, Florida, USA

Jiang Bian

bianjiang@ufl.edu

Department of Health Outcomes and Biomedical
Informatics, University of Florida
Gainesville, Florida, USA

ABSTRACT

Whole genome sequencing (WGS) is quickly becoming the customary means for identification of antimicrobial resistance (AMR) due to its ability to obtain high resolution information about the genes and mechanisms that are causing resistance and driving pathogen mobility. By contrast, traditional phenotypic (antibiogram) testing cannot easily elucidate such information. Yet development of AMR prediction tools from genotype-phenotype data can be biased, since sampling is non-randomized. Sample provenience, period of collection, and species representation can confound the association of genetic traits with AMR. Thus, prediction models can perform poorly on new data with sampling distribution shifts. In this work – under an explicit set of causal assumptions – we evaluate the effectiveness of propensity-based rebalancing and confounding adjustment on AMR prediction using genotype-phenotype AMR data from the Pathosystems Resource Integration Center (PATRIC). We select bacterial genotypes (encoded as k -mer signatures, i.e. DNA fragments of length k), country, year, species, and AMR phenotypes for the tetracycline drug class, preparing test data with recent genomes coming from a single country. We test boosted logistic regression (BLR) and random forests (RF) with/without bias-handling. On 10,936 instances, we find evidence of species, location and year imbalance with respect to the AMR phenotype. The crude versus bias-adjusted change in effect of genetic signatures on AMR varies but only moderately (selecting the top 20,000 out of 40+ million k -mers). The area under the receiver operating characteristic (AUROC) of the RF (0.95) is comparable to that of BLR (0.94) on both out-of-bag samples from bootstrap and the external test ($n=1,085$), where AUROCs do not decrease. We observe a 1%-5% gain in AUROC with bias-handling compared to the sole use of

genetic signatures. In conclusion, we recommend using causally-informed prediction methods for modelling real-world AMR data; however, traditional adjustment or propensity-based methods may not provide advantage in all use cases and further methodological development should be sought.

CCS CONCEPTS

- **Applied computing** → **Computational genomics; Bioinformatics; Computational genomics;** • **Mathematics of computing** → **Causal networks;** • **Computing methodologies** → **Causal reasoning and diagnostics.**

KEYWORDS

antimicrobial resistance, artificial intelligence, biomedical informatics, causal inference, directed acyclic graph, epidemiology, explainability, interpretability, machine learning, propensity score

ACM Reference Format:

Mattia Prosperi, Simone Marini, Christina Boucher, and Jiang Bian. 2021. Assessing putative bias in prediction of anti-microbial resistance from real-world genotyping data under explicit causal assumptions. In *DSHealth '21] Joint KDD 2021 Health Day and 2021 KDD Workshop on Applied Data Science for Healthcare, Aug 14–18, 2021, Virtual*. ACM, New York, NY, USA, 5 pages.

1 INTRODUCTION

Identification of antimicrobial resistance (AMR) from whole genome sequencing (WGS) data is becoming increasingly prevalent due to its ability to render high resolution about the genes and mechanisms that drive resistance and mobility; such resolution cannot be achieved using traditional in vitro antibiotic phenotypic susceptibility testing (antibiogram) [9, 13]. In addition, the development of miniaturized sequencing technology (Oxford Nanopore technology) will drive the emergence of WGS as the de facto method for prediction of AMR genes and resistance phenotype. After the generation of sequence data, computational methods are used to predict the AMR genes and elements. Several of these methods identify the presence of AMR genes using curated databases; for instance, AMR-PlusPlus and the Resistance Gene Identifier align sequence data to the MEGARes and the Comprehensive Antibiotic Resistance Database databases, respectively [1, 6]. Other methods predict the AMR

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DSHealth '21, Aug 14–18, 2021, Virtual

© 2021 Association for Computing Machinery.

phenotype by learning from datasets where bacteria cultured for in vitro antibiotic susceptibility testing are also sequenced, obtaining genotype-phenotype pairs; for instance, interpretable models for multiple antibiotics have been presented [7], using data from the Pathosystems Resource Integration Center (PATRIC) [4].

Development of AMR prediction tools based on genotype-phenotype learning can be biased because the data repositories are made upon non-randomized sampling. Sample provenience and period of collection can affect species representation and confound the association of genetic traits with AMR [19]. In general, any real-world data that is observational in nature can be contaminated with different biases, since the knowledge on the data generation process or the relations among variables involved in such a process is limited [20]. Learning algorithms estimate conditional probability $P(y|X)$ in various forms to predict y from X but the effect of variables in X on to y could be due to unmeasured components Z , which are the true causes of both X and y . Learning models in the presence of bias usually does not affect prediction performance, i.e., $P(y|X)$ can be approximated with negligible error, regardless X being the true causes of y . As a result, prediction models can have high performance in cross-validation, yet perform poorly in real-world settings because of sampling distribution shift, e.g. species prevalence [21].

In this work –under explicit causal assumptions– we evaluate the effectiveness of propensity-based rebalancing and confounding adjustment on AMR prediction with external data characterized by species, time, and location distribution shifts.

2 METHODS

2.1 Potential outcomes and propensity scores

We use Rubin's notation of potential outcomes for calculating exposure or treatment effects [23]. We assume a sample population of N individuals, who were exposed or received a treatment T (binary, for ease) and an health outcome Y was observed after treatment. Each subject i is represented by the tuple $\{T_i, Y_i, X_i\}$, where X is a vector of preexposure variables (e.g., sex, insurance status). We define Y_i^0 and Y_i^1 as the potential outcomes for person i under treatment/exposures $T_i = 0$ and $T_i = 1$, respectively. The individual treatment effect (ITE) $\tau(x)$ is the difference in the average potential outcome for i under both T , conditional on the preexposure vector $X = x$, i.e. $\tau(x) = \mathbb{E}[Y_i^1 - Y_i^0 | X_i = x]$. People cannot have two different exposures (e.g. taking/not taking a therapy or having/not having a genetic trait) at the same time, so only one –factual– outcome can be observed. Thus, the ITE $\tau(x)$ usually cannot be calculated as one outcome is missing. However, if one assumes that potential outcomes are independent of the exposure conditional on preexposure vector, i.e., $\{Y_i^1 Y_i^0\} \perp T | X$, then the ITE can then be calculated under the strongly ignorable treatment assignment (SITA) assumption, as $\tau(x) = \mathbb{E}[Y^1 | T = 1, X = x] - \mathbb{E}[Y^0 | T = 0, X = x] = \mathbb{E}[Y | T = 1, X = x] - \mathbb{E}[Y | T = 0, X = x]$. By assuming SITA and considering the distribution of X , we can calculate the average treatment effect (ATE) τ_{01} as $\tau_{01} = \mathbb{E}[\tau(X)] = \mathbb{E}[Y | T = 1] - \mathbb{E}[Y | T = 0]$. The ATE can also be conveniently expressed as an odds ratio, however, the marginal and conditional effects are not guaranteed to coincide, i.e., the treatment effect measure is not collapsible [5]. ITE and ATE can be calculated with x being equally matched in exposure/control

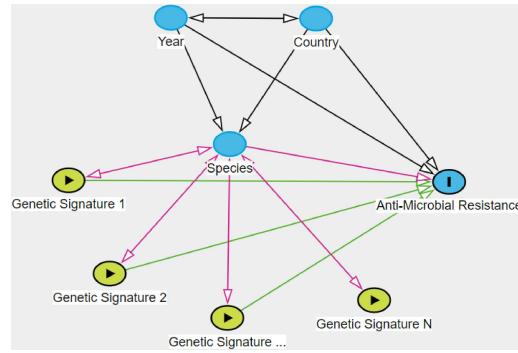


Figure 1: Causal structure (partially directed acyclic graph) for the effect of genetic signatures of bacteria on to antimicrobial resistance, confounded by country of sampling, year, and species.

groups but stratification becomes infeasible as the number of dimensions increases. One way to match individuals without crisp stratification is via propensity scores. The propensity score $\pi(x)$ represents the probability of receiving a treatment or being exposed to $T = 1$ (assuming that the alternative is no treatment or exposure $T = 0$) conditioned on the preexposure covariates X , and is denoted as $\pi(x) = P(T = 1 | X = x)$. Through the conditional probability $\pi(x)$, we can balance the probability of being exposed either to $T = 1$ or $T = 0$ given $X = x$. The propensity score can be estimated using any regression technique from logistic regression to Bayesian additive regression trees [10]. The original dataset can be then rebalanced with respect to the exposure by matching pairs of exposed/unexposed subjects with similar propensity scores, i.e., performing propensity score matching (PSM). Several algorithms can be used for PSM, such as nearest neighbor or caliper [2].

2.2 Proposed approach

As previously described, the effect of a genetic trait of a bacterium on the odds of carrying AMR could be confounded by the fact that the genetic trait itself is characteristic of a species that has been unevenly sampled among antimicrobial-resistant cases. The species could be found to be more prevalent in specific locations or periods of time, and these two variables could affect the overall prevalence of AMR in a population or environment. By considering AMR as an outcome, we can model the genetic trait akin to an exposure confounded by species, locations and times. Figure 1 shows the causal relationships among the aforementioned variables using a partially directed acyclic graph (PDAG). By using d-calculus, e.g. the back-door criterion [17], it is possible to identify a minimal set of variable adjustments for the total effect [18] of one or more genetic signatures on to AMR. Depending on the directions of the arrows between species and genes, the minimum adjustment set would include only the species, or the species, country and year together. Additional arrows between genetic signature pairs or between genetic signatures and the AMR phenotype do not change this adjustment set. Therefore, it is safe to use the species, country and year together as adjustment covariates for effect estimation since they do not add bias for (counterfactual) prediction.

Table 1: Sampling characteristics of the PATRIC genotypes with known tetracycline antibiogram results (n=10,936)

Set (N)	Domain	Variable	Res N (%)	Sus N (%)
Training (9,851)	Country	Canada	131 (59%)	90 (41%)
		Israel	70 (18%)	318 (82%)
		South Africa	281 (56%)	224 (44%)
		UK	567 (43%)	750 (57%)
		USA	2150 (51%)	2059 (49%)
	Year	Other-Unkn.	1278 (40%)	1933 (60%)
		1905-2006	981 (70%)	492 (30%)
		2007-2010	606 (56%)	472 (44%)
		2011-2014	1363 (48%)	1451 (52%)
		2015-	1479 (34%)	2877 (66%)
Test (1,085)	Species	Streptococcus	919 (42%)	1254 (58%)
		Salmonella	1102 (54%)	921 (46%)
		Staphylococcus	322 (16%)	1693 (84%)
		Other	2134 (67%)	1056 (33%)
Test (1,085)	Country/Year Species	USA, 2015-	702 (65%)	383 (35%)
		Streptococcus	1 (100%)	0 (%)
		Salmonella	4 (19%)	17 (81%)
		Staphylococcus	0 (0%)	5 (100%)
		Other	698 (67%)	350 (33%)

Using this explicit causal structure, we define the following data rebalancing and adjusting strategies to handle bias: (i) simple covariate adjustment pooling together all k -mers with country c , year y , and species s variables; (ii) data rebalancing using PSM on c, s, y with respect to AMR phenotype r , i.e. $P(r|c, y, s)$; (iii) covariate ranking and pre-selection based on calculation of ATE of each k -mer adjusted by c, s, y with respect to AMR; (iv) data rebalancing using k -mer clustering (since k -mers are differently distributed within species, this would balance species oversampling regardless the contribution of single k -mers to AMR). For the PSM, we compare three different techniques: nearest neighbor matching with either logistic regression (LR) or random forest (RF) score learning, and coarsened exact matching. The k -mer clustering is done via logistic principal component analysis (LPCA) [12] of the k -mer vectors, then passed on to an agglomerative hierarchical average linkage algorithm, selecting the best number of clusters by maximizing the average silhouette values. Each bias-handling strategy can be applied in conjunction with any machine learning model of choice (whereas the base model is the one that uses only k -mers as input to predict AMR), and here we use boosted LR (BLR) and RF [3, 8].

2.3 Data and experimental settings

We extract genotype-phenotype AMR data and metadata (country, year, species) from the PATRIC repository (<https://www.patricbrc.org>). For this work, we use the tetracycline class of antibiotics as a use case (representing a whole antibiotic class as opposed to single drug compounds, and among the most frequent in the data base). Only bacterial genomes with an associated laboratory test are considered, and only entries that have at least a known metadata attribute. The bacterial genomes (nucleotides) are encoded into k -mers (strings of fixed length k), with $k = 17$, chosen on the basis of optimization from prior studies. Since the total number of k -mers can be overwhelming (hundreds or tens of millions depending on the species diversity), and since k -mers have very long-tailed distributions of frequencies, we select the top 10,000 k -mers based on frequency, and the top 10,000 k -mers based on information

gain with respect to the AMR phenotype (sole training set). These 20,000 are then re-ranked according to each rebalancing strategy, e.g. calculating odds ratio on the basis of PSM for strategy (i) or using the adjusted ATE for strategy (iii), and then the top 5,000 are fed to the internal feature selection of BLR or RF. Countries with low frequency (<90% cumulative, after ordering) are grouped into a single 'other' category. Years are divided into quartiles, and species are grouped at the genus level. A test set is created by selecting instances from a single country, collected in the most recent quartile year. This procedure respects the overlap condition for PSM estimation and yet induces a shifted dataset.

We assess discriminating ability of models using the area under the receiver operating characteristics (AUROC), estimated robustly using out-of-bag samples (10 bootstrap replicates) and on the external shifted dataset. All parameter optimizations (number of boosting iterations and of random trees, between 100 and 6,000) and rebalancing procedures are done internally to the bootstrapping. The PATRIC data preprocessing, LPCA, k -mer counting and extraction –made using KMC3 [11]– are executed on the University of Florida's HiPerGator3 high-performance computing cluster (<https://www.rc.ufl.edu/services/hipergator/>). All subsequent analyses are performed on an Intel laptop with i9-10885H CPU at 2.4GHz and 32MB RAM, using R [22], including libraries cluster, mboost, MatchIt, parallelDist, ranger, ROCR. R scripts and datasets are available at https://github.com/DataIntellSystLab/amr_bias.

3 RESULTS

After querying PATRIC (currently hosting over 80,000 bacterial genomes) according to the inclusion criteria, we obtain 10,936 genotypes and metadata pairs with an antibiogram test result for tetracycline. Of note, additional 8,815 samples meet the inclusion criteria but do not have a legit AMR phenotype. After removing the samples from USA collected in the most recent year quartile, the training

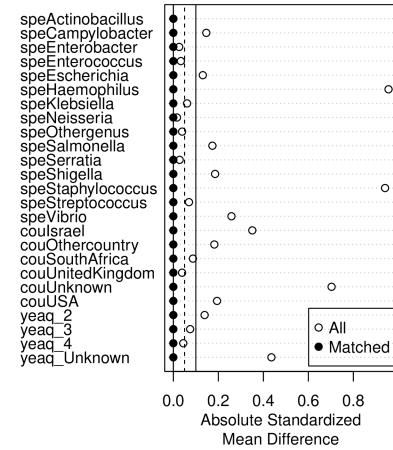


Figure 2: Propensity score matching of sampling variables (spe=species' genus, cou=country, yeaq=quartile year) with respect to AMR phenotype. Black dots show absolute standardized mean difference for the rebalanced data, whilst white dots show the original data.

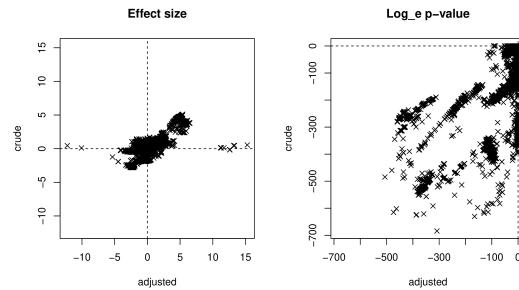


Figure 3: Adjusted vs. crude ATE estimation of each genetic signature (n=20,000) with respect to the AMR phenotype. Left panel shows scatterplot of effect sizes (dashed lines cross zero effect), whilst the right panel compares p-values (dashed lines are p=0.05).

set comprises 9,851 instances, while the external test set includes 1,085. Table 1 shows sampling characteristics of the training and test set, stratified by AMR phenotypes; it is evident how species, country and year differ sensibly among the datasets and strata.

After performing PSM on the AMR phenotypes (with coarsened exact matching being the best method compared to LR or RF nearest neighbor), the sampling variables seem to be perfectly rebalanced with respect to the phenotype (Fig 2). The ATE estimation performed on each k -mer shows that there is dispersion between the crude and the bias-adjusted values, but the direction of the effects does not seem to be strongly affected, with only 8.1% of ATEs changing direction when adjusted (Fig 3). Of note, 20.2% of the k -mers have a p-value > 0.05 when the crude p-value <= 0.05.

The discriminative performance of prediction models on the out-of-bag samples drawn from the training data shows that RF is comparable to BLR using both bias-handling strategies and no bias-handling. The optimal number of boosting iterations for BLR and number of trees for RF stabilizes around 3,000. For BLR, AUROCs range from 0.93 obtained using any of the bias-handling strategies (i)-(iii) to 0.9 without bias handling. For RF, AUROCs range from 0.95 with strategy (iii) to 0.9 without bias handling. The bias-handling strategies yielded an increase in AUROC of 3% to 5% in the out-of-bag validation. When applying the models on the external test dataset, AUROC did not decrease significantly. The gain of the bias-handling strategies was 1%-5%. Fig 4 shows all AUROCs for BLR and RF with respect to each strategy.

4 DISCUSSION

In this work, we explore different strategies to assess putative bias in prediction of AMR from real-world data under explicit causal assumptions. In the PATRIC genotype-phenotype data relative to tetracycline antibiotics, we find strong evidence of sampling and confounding bias with respect to species, country, and year of collection. Nonetheless, the change in effect for each individual genetic signature with respect to ATE when comparing crude and adjusted estimation is mild. Genetic signatures likely already contain information that defines the species, but they likely cannot account for prevalence changes in AMR by country and year. Causally-informed methods improve discriminating ability of AMR as compared to

approaches that use only genetic signatures, but the gain is 5% or less for most configurations.

One limitation of this approach is that we analyze only 20,000 gene signatures out of over 40 million. Half of the pre-selection is based on information gain, which is basically a crude ATE and could be biased. Another possible issue is that the test set is a single realization and thus the AMR distribution shift is non-generalized.

Regardless the use of bias-handling, nonlinear ensemble classifiers like the RF can be computationally burdensome to run even after training. Given that bacterial sequencing is being increasingly done in real-time and mobile settings [15], development of fast classifiers is warranted, thus the choice of BLR seems favorable.

Here, we restrict interest to WGS but extension to metagenomics data warrants future investigations. Studying bias in this setting is complicated by the prevalence of multiple – sometimes putative – AMR pathogens and species, which will confound the k -mer composition. Moreover, geographical location has shown to have a substantial effect of metagenomes [14, 16] so we hypothesize that our results will be more pronounced in a metagenomics setting.

In conclusion, we recommend the usage of causally-informed methods for the development of computational AMR prediction models using observational, real-world data, when there is evidence of bias in species sampling; however, traditional adjustment or propensity-based methods may not provide advantage in all use cases and further methodological development should be sought.

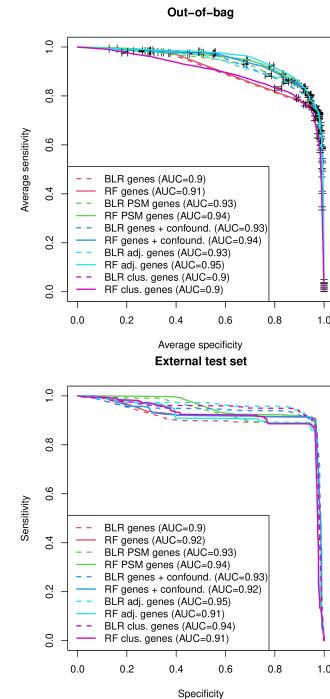


Figure 4: Discriminating ability in prediction of AMR phenotype by random forest and boosted logistic regression with different bias-handling strategies (vs. none) on ten out-of-bag samples (top panel) and external test set (bottom panel).

ACKNOWLEDGMENTS

This work was in part supported by US grants NIH NIAID R01AI145552, NIH NIAID R01AI141810, and NSF SCH 2013998.

REFERENCES

- [1] Brian P Alcock, Amogelang R Raphenya, Tammy T Y Lau, Kara K Tsang, Mégane Bouchard, Arman Edalatmand, William Huynh, Anna-Lisa V Nguyen, Annie A Cheng, Sihan Liu, Sally Y Min, Anatoly Miroshnichenko, Hiu-Ki Tran, Rafik E Werfalli, Jalees A Nasir, Martins Oloni, David J Speicher, Alexandra Florescu, Bhavya Singh, Mateusz Faltyń, Anastasia Hernandez-Koutoucheva, Arjun N Sharma, Emily Bordeleau, Andrew C Pawlowski, Haley L Zubyk, Damion Dooley, Emma Griffiths, Finlay Maguire, Geoff L Winsor, Robert G Beiko, Fiona S L Brinkman, William W L Hsiao, Gary V Domselaar, and Andrew G McArthur. 2019. CARD 2020: antibiotic resistance surveillance with the comprehensive antibiotic resistance database. *Nucleic Acids Research* 48, D1 (10 2019), D517–D525.
- [2] Peter C Austin. 2014. A comparison of 12 algorithms for matching on the propensity score. *Statistics in Medicine* 33 (2014), 1057–1069.
- [3] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32.
- [4] James J Davis, Alice R Wattam, Ramy K Aziz, Thomas Brettin, Ralph Butler, Rory M Butler, Philippe Chlenski, Neal Conrad, Allan Dickerman, Emily M Dietrich, Joseph L Gabbard, Svetlana Gerdes, Andrew Guard, Ronald W Kenyon, Dustin Machi, Chunhong Mao, Dan Murphy-Olson, Marcus Nguyen, Eric K Nordberg, Gary J Olsen, Robert D Olson, Jamie C Overbeek, Ross Overbeek, Bruce Parrello, Gordon D Pusch, Maulik Shukla, Chris Thomas, Margo VanOeffelen, Veronika Vonstein, Andrew S Warren, Fangfang Xia, Dawen Xie, Hyunseung Yoo, and Rick Stevens. 2019. The PATRIC Bioinformatics Resource Center: expanding data and analysis capabilities. *Nucleic Acids Research* 48, D1 (10 2019), D606–D612.
- [5] Vanessa Didelez and Mats Julius Stensrud. 2021. On the logic of collapsibility for causal effect measures. *Biometrical Journal* (2021), 1–8.
- [6] Enrique Doster, Steven M Lakin, Christopher J Dean, Cory Wolfe, Jared G Young, Christina Boucher, Keith E Belk, Noelle R Noyes, and Paul S Morley. 2019. MEGARes 2.0: a database for classification of antimicrobial drug, biocide and metal resistance determinants in metagenomic sequence data. *Nucleic Acids Research* 48, D1 (11 2019), D561–D569.
- [7] Alexandre Drouin, Gaël Letarte, Frédéric Raymond, Mario Marchand, Jacques Corbeil, and François Laviolette. 2019. Interpretable genotype-to-phenotype classifiers with performance guarantees. *Scientific Reports* 9, 1 (mar 2019), 4071.
- [8] J. Friedman, T. Hastie, and R. Tibshirani. 2000. Additive Logistic Regression: a Statistical View of Boosting. *The Annals of Statistics* 38, 2 (2000).
- [9] Rene S. Hendriksen, Valeria Bortololia, Heather Tate, Gregory H. Tyson, Frank M. Aarestrup, and Patrick F. McDermott. 2019. Using Genomics to Track Global Antimicrobial Resistance. *Frontiers in Public Health* 7 (2019), 242.
- [10] Jennifer L. Hill. 2011. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics* 20, 1 (2011), 217–240.
- [11] Marek Kokot, Maciej Dlugosz, and Sebastian Deorowicz. 2017. KMC 3: counting and manipulating k-mer statistics. *Bioinform.* 33, 17 (2017), 2759–2761.
- [12] Andrew J Landgraf and Yoonkyung Lee. 2020. Dimensionality reduction for binary data through the projection of natural parameters. *Journal of Multivariate Analysis* 180 (2020), 104668.
- [13] Sarah Lepuschitz, Thomas Weinmaier, Katharina Mrazek, Stephan Beisken, Johannes Weinberger, and Andreas E Posch. 2020. Analytical Performance Validation of Next-Generation Sequencing Based Clinical Microbiology Assays Using a K-mer Analysis Workflow. *Frontiers in Microbiology* 11 (2020), 1883.
- [14] Tanya M. Monaghan, Tim J. Sloan, Stephen R. Stockdale, Adam M. Blanchard, Richard D. Emes, Mark Wilcox, Rima Biswas, Rupam Nashine, Sonali Manke, Jinal Gandhi, Pratishtha Jain, Shrejal Bhotmange, Shrikant Ambalkar, Ashish Satav, Lorraine A. Draper, Colin Hill, and Rajpal Singh Kashyap. 2020. Metagenomics reveals impact of geography and acute diarrheal disease on the Central Indian human gut microbiome. *Gut Microbes* 12, 1 (2020), 1752605.
- [15] Marco Oliva, Franco Milicchio, Kaden King, Grace Benson, Christina Boucher, and Mattia Prosperi. 2020. Portable Nanopore Analytics: Are We There Yet? *Bioinformatics* (2020).
- [16] Edoardo Pasolli, Francesco Asnicar, Serena Manara, Moreno Zolfo, Nicolai Karcher, Federica Armanini, Francesco Beghini, Paolo Manghi, Adrian Tett, Paolo Ghensi, Maria Carmen Collado, Benjamin L. Rice, Casey DuLong, Xochitl C. Morgan, Christopher D. Golden, Christopher Quince, Curtis Huttenhower, and Nicola Segata. 2019. Extensive Unexplored Human Microbiome Diversity Revealed by Over 150,000 Genomes from Metagenomes Spanning Age, Geography, and Lifestyle. *Cell* 176, 3 (2019), 649–662.e20.
- [17] J. Pearl, M. Glymour, and N.P. Jewell. 2016. *Causal Inference in Statistics: A Primer*. Wiley.
- [18] Emilia Perković, Johannes Textor, Markus Kalisch, and Marloes H. Maathuis. 2015. A Complete Generalized Adjustment Criterion. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence* (Amsterdam, Netherlands) (UAI'15). AUAI Press, Arlington, Virginia, USA, 682–691.
- [19] Mattia Prosperi, Taj Azarian, Judith A. Johnson, Marco Salemi, Franco Milicchio, and Marco Oliva. 2019. Unexpected Predictors of Antibiotic Resistance in Housekeeping Genes of *Staphylococcus Aureus*. In *Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics* (Niagara Falls, NY, USA) (BCB '19). Association for Computing Machinery, New York, NY, USA, 259–268.
- [20] Mattia Prosperi, Yi Guo, Matt Sperrin, James S. Koopman, Jae S. Min, Xing He, Shannan Rich, Mo Wang, Iain E. Buchan, and Jiang Bian. 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence* 2, 7 (July 2020), 369–375.
- [21] Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. 2009. *Dataset Shift in Machine Learning*. The MIT Press.
- [22] R Core Team. 2020. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- [23] D.B. Rubin. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66, 5 (1974), 688–701.