How accurately can machine learning algorithms predict a person's future?

Emily Cantrell, Princeton University (presenting author; e.cantrell@princeton.edu)
Pranay Anchuri, Princeton University
Hanzhang Ren, Stanford University
Matthew Salganik, Princeton University

Applications of artificial intelligence (AI) are rapidly increasing in public policy. One such application is the use of machine learning algorithms in public services to make predictions about individual people's futures. For example, several U.S. child protective services agencies use algorithms to predict the likelihood that a child will be removed from the home due to maltreatment.[i] Predictive algorithms are also being used or piloted in criminal justice,[ii] homelessness prevention,[iii] medical services,[iv] school dropout prevention,[v] and more.[vi] To understand the appropriateness of using such algorithms, we need to understand how accurately a person's future can be predicted.

We investigate the predictability of hundreds of life outcomes for youth and their families, and identify patterns in what types of outcomes are most (and least) predictable. Specifically, we use in-depth survey and observational data collected from a child's birth to age 9 in a U.S. longitudinal study[vii] to predict approximately 500 outcomes at the child's age 15. This dataset includes thousands of variables about children and their families in a wide range of life domains. We test dozens of machine learning approaches for every outcome to ensure a reasonable estimate of the best possible predictive accuracy that can be achieved with these data. Our main finding is that most outcomes in our set are not very predictable. For approximately half of the outcomes, our algorithms do not make predictions any better than a null model that predicts everyone's outcome to be equal to the mean of the training data. However, some outcomes are much more predictable than others, and we identify patterns in what types of outcomes are most predictable.

These findings call for caution about predictive accuracy when implementing predictive algorithms in public services. However, they also raise questions about how sample size and other characteristics of data affect predictive accuracy. In future work, we plan to explore questions about the capabilities and limitations of predictive algorithms with Dutch administrative register data, which has similar characteristics to the administrative register data of Nordic countries.

---

[i] Samant, A., Horowitz, A., Xu, K., & Beiers, S. (2021). Family surveillance by algorithm: The rapidly spreading tools few have heard of. American Civil Liberties Union (ACLU).
[ii] APPR. (2023). About the Public Safety Assessment: Public Safety Assessment Sites. https://advancingpretrial.org/psa/psa-map/
[iii] Shinn, M., Greer, A. L., Bainbridge, J., Kwon, J., & Zuiderveen, S. (2013). Efficient targeting of homelessness prevention services for families. *American Journal of Public Health*, *103*(S2), S324-S330.
[iv] Pan, I., Nolan, L. B., Brown, R. R., Khan, R., van der Boor, P., Harris, D. G., & Ghani, R. (2017). Machine learning for social services: a study of prenatal case management in Illinois. *American Journal of Public Health*, *107*(6), 938-944.
[v] Wang, A., Kapoor, S., Barocas, S., & Narayanan, A. (2023). Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. Available at SSRN.
[vi] Levy, K., Chasalow, K. E., & Riley, S. (2021). Algorithms and Decision-Making in the Public Sector. *Annual Review of Law and Social Science, 17.*
[vii] Future of Families and Child Wellbeing Study. https://ffcws.princeton.edu/about