Child Welfare Fact Sheet

Predictive Risk Modeling for Child Protection

Every year, more than 3.6 million referrals involving about 6 million children are made to child protective services (CPS) in the United States.[1] It is estimated that around one in three American children will have been investigated by CPS for maltreatment by age 18.[2] Flooded with calls, and short on resources, CPS systems are struggling to identify and protect children at risk. Decision-support tools such as predictive risk models can help CPS workers navigate an environment where it is hard to separate "noise" from "signal."

How CPS decisions are made

CPS decision making generally relies on human judgment—with some additional support from actuarial tools like Structured Decision Making®. Actuarial tools generally require workers to answer a list of questions and enter additional data. These answers are weighted, and the tool produces a score. The weights are universal and are most often not validated against local data. These tools support decision making by frontline staff but do not lead to required decision paths.

What is a predictive risk model?

Predictive risk models (PRMs) are an improvement over actuarial tools in the way they are built, the data they use, and their accuracy. Specifically, they are trained[3] using local data, are able to draw upon a much broader set of information fields, and have been shown to have better classification accuracy than earlier generation actuarial tools.

PRM tools can be used at various decision points in the child welfare system—but they all work in essentially the same way. When the tool is triggered, such as when a hotline referral is made, the PRM tool automatically harvests data fields about each individual on the call. These fields include previous interactions recorded in the Statewide Automated Child Welfare Information System (SACWIS) or new Comprehensive Child Welfare Information System (CCWIS) or other administrative data. Each of these fields is combined to generate a score instantaneously. Because PRM scores rely on data in administrative systems, they require no additional data entry.

Data used by PRM

Accurate PRM tools can be built using only data from the SACWIS/CCWIS systems.[4] However, for jurisdictions with integrated data systems, criminal justice or behavioral health fields can be added if it increases accuracy.

Generating a score through combined data fields

When a PRM is triggered, each data field is combined by a set of rules (called an algorithm) to generate a score instantaneously. The algorithm depends on the specific machine-learning method that has been used to train the model. If the tools are based on regression methods, then weights are applied to different attributes. The training process looks at how well each attribute is able to predict a future outcome (such as chronic future involvement in CPS or placement). It is important to be aware that the relationships that emerge do not necessarily reflect any causal relationship between attributes and the future outcome.

The PRM score is never used to make decisions. It is treated as a decision-support tool and an additional piece of information for workers. Constant monitoring and evaluation by leadership is necessary to ensure that staff are making appropriate use of the PRM score as well as their own judgment, and not overly relying on the tool.

Child protection use cases

A PRM tool that **generates a risk score at referral** could be useful in many ways, including the following:

1. **Process oversight:** To help understand and improve practice.

   Agencies can use a data-driven tool to estimate the risk of adverse outcomes (that is, placement in foster care) and as a so-called yardstick to measure current processes, identifying elements that are working well and areas for improvement.

2. **Hotline screening:** To help make screening decisions more consistent, equitable, and accurate.

   Rather than simply relying on the partial information communicated on calls alleging abuse or neglect and hoping that hotline screeners fully review and properly weight the historical information they have access to, PRM provides a standardized manner for incorporating data into the call-screening process.

3. **Enhanced supervision:** To support clinical supervision via triaging of investigations and open cases.

   Child protection agencies could establish a threshold (for example, highest risk 5 percent or 10 percent for flagging complex maltreatment reports assigned for investigation. These flags would accompany the information transmitted to supervisors overseeing those investigations and could be used to support supervisors to prioritize time spent consulting on investigations, adjust investigatory workloads, confirm the completeness of investigations, and review screening decisions before signing off on service decisions.

4. **Family First candidacy:** To help classify and prioritize candidates for federally funded prevention services.

   The Family First Prevention Services Act provides states with new optional Title IV-E[5] funds to provide families with children at risk of foster care placement up to 12 months of mental health services, substance abuse treatment, and in-home skill-based parenting training and family therapy services.[6] To receive these prevention services under the federal reimbursement scheme, a child must be defined as a candidate for foster care, or as a child identified as being at imminent risk of entering foster care but who can remain safely at home so long as preventative services are provided. Agencies could use a PRM that classifies children based on the risk of future foster care placement as one of several approaches to establishing candidacy.

Other places where a PRM score could be useful include the following:

1. Generating a risk score at the start of a removal to help identify children at low risk of reunification in order to expedite to permanency.

2. Generating a risk score at reunification to help identify children who are at risk of reentry into foster care to prioritize support services.

3. Regular risk-scoring of all children with case openings to identify children at greatest risk of adverse outcomes to help focus supervisory efforts.

Why consider predictive risk models?

✓ High stakes decisions made under time pressure require accurate decision-support tools that are easy to use.

✓ PRMs generally perform better than the currently used actuarial tools because they are bespoke for a jurisdiction, do not rely on operator-entered data, and can draw on a wide range of cross-sector information.

✓ With heightened awareness of abuse and the opioid crisis,[7] national referral rates have been steadily rising from 2013 to 2017, requiring better support for agencies.

✓ Many jurisdictions have increasingly rich and useful collections of administrative data, including data from other programs such as education, health care, and other social services.

✓ PRMs can be an additional resource at key decision-making points because frontline workers struggle to weight complex factors, even when they have access to rich data.

✓ PRMs allow high quality information about cross-sector service interactions to inform our understanding of risk and protective factors.

✓ An independent evaluation has shown that a call screening PRM implemented in Allegheny County, Pennsylvania, improved the efficiency of screening decisions and reduced racial disparities.

Addressing PRM concerns

A frequent concern is that **PRMs predict system outcomes or system responses** (like placement), so might not tell us about true underlying risk of serious harm to the child. However, we have found that PRM tools trained to identify children at risk of removals are also highly sensitive to identifying other risks such as maltreatment-related fatalities and other objective measures of harm. Specifically, our work in California suggests that among children who experience a near-fatality or fatality because of maltreatment, nearly 60 percent would have been flagged at an earlier referral as falling in the top 10 percent of risk had a PRM been implemented.

People often worry that data used for a **PRM might be racially biased**. This concern is valid, given the acknowledged bias in contributing data sets; the question of racial bias should be directly explored. This concern is best addressed by studying how well the PRM performs for different racial subgroups—that is, does the PRM gauge risk similarly for different racial groups? It is also important to keep in mind that the current approach to decision making in child protection is not free of bias. PRMs can be a useful tool to identify bias in existing decision making processes by documenting where there is unwarranted variation in the system's response to families based on race, rather than risk.

Consent and privacy concerns—around the data used for a PRM and the screening score that is generated—are understandable. For this reason, we prefer to use data that a given jurisdiction has already determined should be used in the existing decision-making process. In addition to having the legal rights to use data for a PRM, an agency should make efforts to achieve social license for this new use—open consultation with community and advocacy groups' concerns with privacy can be very helpful. When it comes to the screening scores generated by a PRM, the agency needs to make thoughtful decisions on who will have access to those scores and why. These policy decisions should take account of ethical concerns and should be clearly communicated to workers and the community.

Guardrails: Why and what?

As part of each new project, we define a set of guardrails that will promote trusted use of PRMs.

AGENCY LEADERSHIP

Implementing a PRM tool requires multiple choices and trade-offs. It is important that the agency (working closely with the vendor or research partner) makes these key policy and practice choices and communicates these to its community. Ownership of the tools, code, and all other material should also lie with the agency and not the vendor.

TRANSPARENCY AND ETHICS

The vendor of the PRM tool should provide documentation to the agency—and information about use of the PRM tool should be available to the public. The document should allow an assessment of how the PRM tool was trained (data fields used, and so on), its accuracy (including for subpopulations), and the specific use-cases for which the tool was built.

If there are concerns about the use of PRMs, then an independent ethical evaluation of the tool could also be valuable. Such a report could set out the potential value and possible pitfalls of the proposed use-case and provide suggestions as to how to manage these.

COMMUNITY VOICE

Because of the novelty of PRMs, and the fact that they use data from the community to train the model, the agency needs to initially obtain social license from the community to use their data in this way. Seeking out community members' input—especially from those who are most likely to be affected by the tool—before, during, and after implementation is an important component of obtaining that license.

EVALUATION AND MONITORING

The impact of the PRM tool on decisions needs to be carefully evaluated. Commissioning an initial independent impact evaluation can be valuable to establish whether the promoted benefits have been realized. However, it is crucial that leaders engage in continuous monitoring of PRM tools to ensure that they maintain their accuracy and utility.

How PRM is being used in child protection

Allegheny Family Screening Tool. Initial research developed by Vaithianathan and Putnam-Hornstein in 2016 led to building and implementing the Allegheny Family Screening Tool, the first-ever use of an automated risk algorithm to support the screening of child maltreatment allegations. This work in Allegheny County, Pennsylvania, established that PRM scores could be used by child protection hotline staff to triage children based on the likelihood of future involvement with the child protection system. An independent evaluation conducted by researchers at Stanford University concluded that the use of the tool was not associated with any harmful effects on children or families, improved the accurate identification of children in need of services, had no detectable effect on decisions to screen out children without investigation, and was associated with a modest but detectable reduction in racial disparities in case openings.
Read more: www.alleghenycountyanalytics.us

Douglas County Decision Aide. Vaithianathan and Putnam-Hornstein led a feasibility study in 2017 and went on to implement the Douglas County Decision Aide (DCDA), a hotline screening tool, in Douglas County, Colorado, in early 2019. Two aspects of the DCDA are notable: it is a PRM built using only child welfare data and public benefit eligibility data, and it will be evaluated via a randomized controlled trial, due to be completed in February 2020 by an independent team at Cornell University.

California Proof of Concept PRM. In 2016, Vaithianathan and Putnam-Hornstein began PRM work in California, modeled closely on the Allegheny Family Screening Tool (AFST). Given that the state does not have an integrated data system, the goal of this project was to assess whether a PRM built exclusively from child protection records could approach the accuracy of the AFST. This proof of concept also included

comparisons between scores generated through a PRM and risk levels assigned through the use of the SDM® Family Risk Assessment tool used in California. Findings indicated that the PRM was more accurate in identifying children who would have chronic or intensive involvement with the child protection system. The state is currently developing implementation plans with county stakeholders. Read more: www.datanetwork.org

Endnotes

[1] U.S. Department of Health and Human Services. "Child Maltreatment 2014." Washington, DC: U.S. Department of Health and Human Services, Administration for Children and Families, Administration on Children, Youth and Families, Children's Bureau, 2016. Available at http://www.acf.hhs.gov/programs/cb/research-data-technology/statistics-research/child-maltreatment.

[2] Kim, H., C. Wildeman, M. Jonson-Reid, and B. Drake. "Lifetime Prevalence of Investigating Child Maltreatment Among U.S. Children." *American Journal of Public Health*, vol. 107, no. 2, 2017, pp. 274–280.

[3] Training a model refers to the statistical process of mining data to look for patterns that can then be embedded into the predictive risk model.

[4] In our experience, a small number of calls (around 10 percent) will involve situations where there is no information in the administrative systems about any of the people involved in the incident. In this case, the processes that were used before the adoption of the PRM tool will be used.

[5] Funds derived from Title IV-E of the Social Security Act.

[6] Administration for Children and Families. "Information Memorandum NEW LEGISLATION—Public Law 115-123, the Family First Prevention Services Act within Division E, Title VII of the Bipartisan Budget Act of 2018." Memorandum to state, tribal, and territorial agencies administering or supervising the administration of Title IV-E and/or Title IV-B of the Social Security Act, April 12, 2018. https://www.acf.hhs.gov/sites/default/files/cb/im1802.pdf.

[7] https://aspe.hhs.gov/pdf-report/relationship-between-substance-use-indicators-and-child-welfare-caseloads.

More Information

This summary has been prepared by a partnership focused on the safe and ethical use of predictive analytics for child protection, led by the following: Professor of Health Economics and Social Data Analytics and Director of the Centre for Social Data Analytics, **Professor Rhema Vaithianathan**, Auckland University of Technology (New Zealand) and University of Queensland (Australia); Associate Professor of Social Work and Co-Director of the Children's Data Network, **Emily Putnam-Hornstein**, University of Southern California (Los Angeles); Public Policy and Human Services Specialist and Vice President and Director of Human Services at Mathematica **Matthew Stagner**; Senior Researcher and Lead for State and Local Child Welfare at Mathematica, **Elizabeth Weigensberg**.