Formative Modeling of Foster Care Work: A Cognitive Work Analysis Approach

Connor Wurst, Huei-Yen Winnie Chen, and Kenneth Joseph
University at Buffalo, State University of New York
Amherst, New York

ABSTRACT

In this paper we present the promise of the Cognitive Work Analysis (CWA) methodology, particularly abstraction hierarchy modeling, in the foster care domain. There is increasing interest in applying machine learning decision aids to foster care decision making, but that interest is accompanied by concerns that those aids may perpetuate systemic bias or be largely context-blind. Modeling the work conducted at different levels of the domain offers unique insights into where bias may enter the system as well as possible design implications for these future decision aids. This project models two major areas of work in the domain, management of individual cases and management of overall programs offered. These work areas are then considered in the first 3 levels of the abstraction hierarchy to display the promise that this model can hold for the domain in future work, particularly when supported with more naturalistic studies.

INTRODUCTION

Foster care is a complex system in which youth, family, and government priorities must be balanced by a third-party agency which is frequently underfunded and understaffed. This complexity means that new solutions to resource allocation and other decisions in the domain are much needed, and many solutions are proposed from the fields of machine learning and data science. Predictive models based on historical and administrative data are frequently used for risk assessment, however there are important questions surrounding how to balance these with traditional casework (Boyas, 2017). Additionally, both machine learning and traditional casework are frequently affected by systemic and data-based biases which may arise from a number of different sources (Suresh & Guttag, 2019). In this context, biases occur when data or personal experiences are harmfully misused to perpetuate inequalities. To better understand the role of bias in these decisions and how machine learning may efficiently and effectively fit into the decision-making processes of social workers in the foster care setting, a rigorous model of foster care work and decisions must be established.

This work is a preliminary exercise in a larger interdisciplinary collaboration between social work, computer science, and industrial engineering departments of the University at Buffalo and a local foster care agency. The project at large aims to create a better understanding of the role of bias in the foster care system as well as deliver a context-informed decision aid. The proposed model of work would have extensive implications within those goals, particularly with respect to understanding bias and designing future decision aids.

BACKGROUND

Decision Making in Social Work

The funding and manpower constraints on foster care settings have resulted in extensive research that aims to improve their best practices. Component analysis of social worker survey results show that there is a complex interplay between internal and external references when social workers make child removal decisions (Dettlaff, Christopher Graham, Holzman, Baumann, & Fluke, 2015), offering one potential source of bias. In reunification decisions, child wellness outcomes varied with a wide variety of individual and family factors (Biehal, Sinclair, & Wade, 2015), which highlights the complex dataset which must inform decisions in the domain. While risk assessment tools have been developed to assist with initial screening calls, later analysis revealed social workers rarely utilized the recommendation (Chouldechova et al., 2018). Similarly, social worker decision making was shown to be far more heuristic than analytical in think-aloud studies (Zeijlmans, López López, Grietens, & Knorth, 2019). These prior studies suggest room for improvements in analytical decision support tools, given that decisions made in the foster care setting are complex and involve both case data and interpretation of that data given prior experience.

Bias has been identified to play a large role in social work decision making, particularly in the process of selecting what information will be used to make the decision (Munro, 1999)., but including bias in a model of foster care work would allow for a better understanding of how it initially occurs and is perpetuated. This includes the bias that may be introduced by way of the data used to train machine learning methods for decision aids and the presentation of the decision aid as well. A formative model which correctly captures the overall goals and constraints of the work domain would not only describe how work is currently done, but how it might be done more efficiently and ethically.

Work Domain Analysis

Cognitive Work Analysis (Vicente, 1999) was born out of the control needs of the nuclear energy field, and grows increasingly important as workplaces continue to grow more complex. The aim of the framework was to model complex sociotechnical domains in such a way that flexible information systems that supported human cognition could be developed. The approach begins at an ecological level and continues to drill down into the specifics of the workplace while considering the constraints placed on worker action at each level. By considering the constraints imposed on the work rather than current practices, CWA functions as a formative approach to work analysis that can provide new and flexible solutions (Vicente, 1999).

CWA has been applied extensively in a wide range of fields. Informatics, error prevention, and decision support have been extensively researched through the CWA lens in the healthcare domain (Jiancaro, Jamieson, & Mihailidis, 2014). While some have claimed that a gap still exists between CWA and concrete design work, design toolkits leveraging the framework have been developed (Read, Salmon, Lenné, & Stanton, 2015). Additionally, design projects including interface design, function allocation, and team design in industrial and constant process domains have all been supported by CWA insight (Read, Salmon, & Lenné, 2012). The complex and dynamic information needs of aviation have been a rich field of CWA study as well, including accident analysis (Coury & Schulze, 2004) and communications planning (McIlroy & Stanton, 2011).

The first and most ecological level of CWA is Work Domain Analysis (WDA), largely represented by an Abstraction Hierarchy, which breaks the system down both in means-end and part-whole hierarchies (Rasmussen et al., 1994). This method focuses on identifying the contingencies between the highest level, most abstract goals of the organization and the concrete physical objects and tasks that drive them. By identifying these contingencies, a rigorous understanding of system structure can be developed allowing practitioners to build better tool aids and information systems. Additionally, communicating information to domain experts in ways that leverage abstractions common to their expertise can be a powerful aid in interface design.

The abstraction hierarchy has continued to evolve and be broadly applied following Rasmussen's original framing. It's been used extensively in data visualization methodologies (Rouse, Pennock, Oghbaie, & Liu, 2017), particularly in Ecological Interface Design (Vicente, 2002). Other extensions to the AH methodology include considerations of team performance (Stanton et al., 2004), inclusion of cognitive as well as physical objects (Carden, Goode, & Salmon, 2019), and most crucially for this work, application of the method to intentional domains (Wong, Sallis, & O'Hare, 1998).

Intentional domains vary from the early CWA applications (such as nuclear power) in that they are constrained less by physical processes and natural law; they are instead spaces in which human decision and cognition drives much of the work. Wong has contributed a number of publications on the application of CWA to intentional domains including ambulance dispatch systems (Wong et al., 1998) and analysis of the VALCRI algorithm used in criminal justice (Paudyal & Wong, 2018).

This work aims to present the promise of the AH to the foster care domain, as far more intensive naturalistic study is needed in order to fully realize these promises. An initial formulation of the AH is presented to explain the methodology as well as shed light on the complexity of this sociotechnical system, but interviews with subject matter experts and review of workplace documentation will be necessary to finalize the model.

RESEARCH OBJECTIVES

While the impact of machine learning in foster care has been extensively researched and called for by "evidence based" legislation such as the Family First bill ("116th Congress Bill Profile H.R. 2702, Family First Transition and Support Act of 2019, Bill Profile," 2019), few agencies have actually implemented these tools. This offers a greenfield opportunity which strongly fits the capabilities of WDA as a formative method. WDA was designed specifically not to describe work as done, rather by capturing the constraints of the domain it is able to offer insight into new design and possibilities. Careful analysis of the work domain in light of both program and case management will allow for future development of a decision-making aid which will make the domain a much more technical sociotechnical system.

The VALCRI CWA (Paudyal & Wong, 2018) work offers a number of similarities to this project; both focus on intentional domains with complex sociotechnical implications and are interested in the application of algorithmic decision aids in those systems. However, the VALCRI analysis was focused on increasing the transparency of an algorithm's functioning after its implication, whereas this project aims to better understand the work domain before the algorithm design begins. Of particular focus will be how the biases highlighted in the foster care decision making literature can be captured in the CWA framework. A more robust understanding of the origin and upkeep of biases will allow for the development of decision aids which can mitigate bias in the domain.

A final point of interest for the foster care domain arises when considering the part-whole breakdown of the abstraction hierarchy. Rather than a clear delineation of components and subsystems, the domain is best broken down into program management and case management. Program management consists of deciding what large-scale offerings the foster agency will present to the community and focuses largely on resource management, whereas case management is the social workers' handling of individual cases and youth outcomes. Program management often constrains case decisions, and the dependencies between the two levels was heavily emphasized by subject matter experts.

Given the current state of the domain and the previous lack of WDA applications in it, two main research objectives arise. First, how can the conflicting aims of program management and case management be properly captured in an abstraction
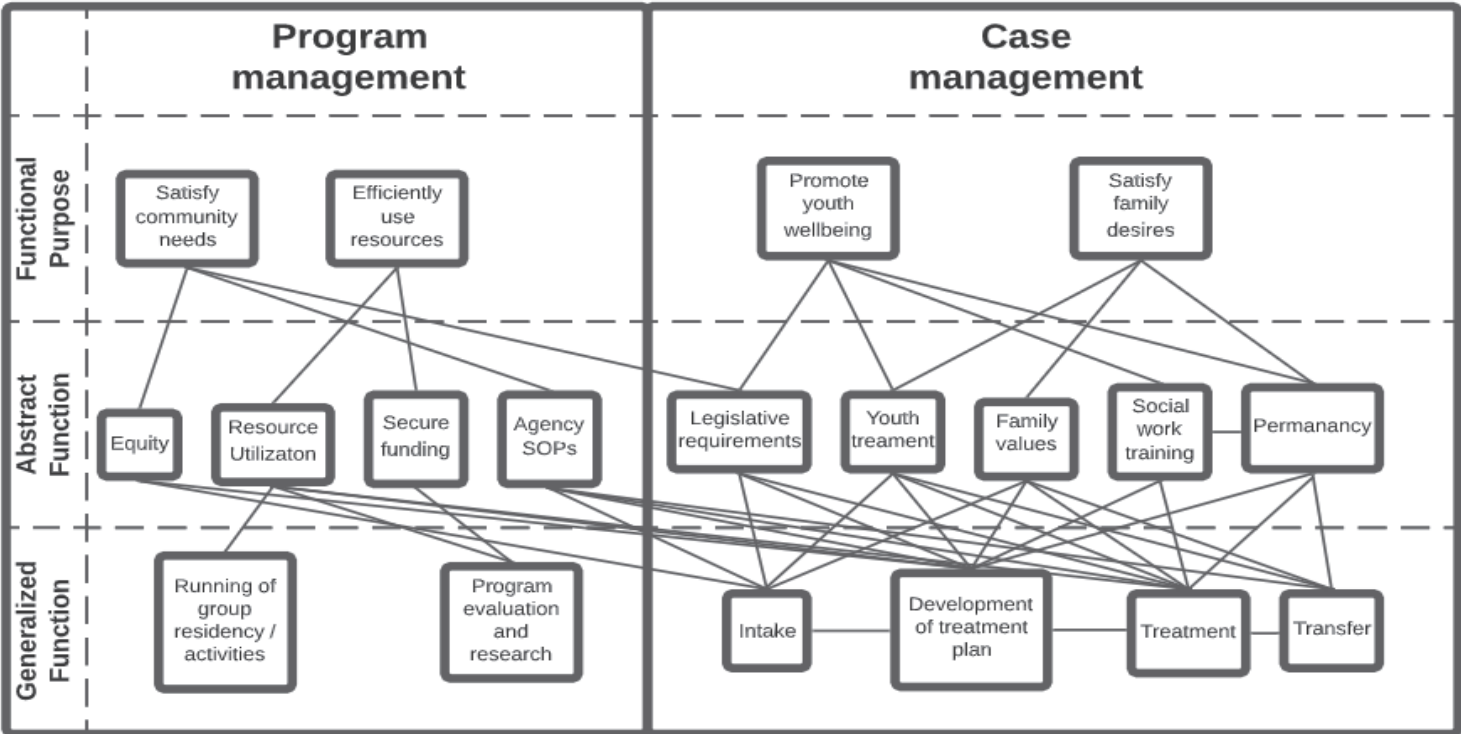
Fig. 1: Partial Abstraction Hierarchy of Foster Care Work

hierarchy? Second, how might systemic bias be represented in that hierarchy, and what insight might be gained from that representation?

MODEL AND DISCUSSION

The current abstraction hierarchy is focused on the top 3 levels, which capture the overall goals of the organization (functional purpose), the values and priorities needed to meet those goals (abstract function), and the activities needed to be completed to optimize the values (generalized functions). Typically, two final levels (physical function and physical form) are included to represent the most concrete tasks and items in the system, but these are left for future work. Items in each of the levels are connected by means-ends links, which indicate dependencies. This model was developed following informal conversation with collaborators well-versed in the child welfare system as well as review of various documentation used in the domain. This partial hierarchy is being used to highlight some of the initial complexities of the system, while the addition of the most concrete levels (physical function and form) will be used to drive the eventual creation of any decision-making aids.

Program Management

As previously mentioned, the part-whole hierarchy consists of a program management level in which social workers decide what services to offer and how they will be run as well as a case management level in which individual treatment plans are handled. A major focus of this level is the matching of programs to community needs, which may include changing the offerings of the agency. However, those needs

need to be balanced with the other functional purpose of efficiently using resources. This creates the first insight into frictions within the system, particularly if meeting community needs is resource intensive. The efficient use of resources is also impacted by what resources the agency has to utilize, adding extra weight to the abstract function of securing funding. There are generally extensive requirements to be met by agencies to receive various state and federal funding, which further constrains choices in both program and case management.

Equity is found at the abstract function level, indicating it as a major value of the system, and its implementation is key to any discussion of bias within the system. If social workers' internal or systemic biases are a large factor of treatment plan decisions, equity will fail to be met. This then results in a failure of the functional purpose to satisfy community needs. Given that the treatment plan development also drives youth treatment and eventual outcomes, this area should be central to future investigation and will be discussed further in the case management section.

At the generalized function level, program management breaks down into two major tasks to be accomplished: running the programs and ensuring that the programs continue to function well by way of research and evaluation. Program operation is another point that was emphasized by subject matter experts as particularly constraining. When deciding what programs a child will be placed in, the social workers need to consider not just how the program will benefit them, but also how that youth will fit into the current group in that program. Graphically, this insight can be seen by observing that both development of treatment plan and running of group activities impact the resource utilization abstract function.

This is of particular importance given how many of the programs are residential in nature.

Case Management

In the realm of managing individual cases, social workers seem to follow two functional purposes: promoting youth wellbeing and satisfying the family's needs or desires for the case. While in most cases these abstract goals can easily coexist, more difficult cases may cause friction between them. Disagreement in regards to a youth's best interests amongst case workers and family members will place these two goals at odds, increasing the complexity of the interactions amongst the less abstract levels.

At the abstract function level, family preferences are still of strong consideration. However, the norms of social work training may differ from those values, and these norms may very well include biases that significantly affect caseworker decisions. Past studies have shown that bias often still affects clinical decisions and that there is a need for more extensive bias mitigation in training (Featherston et al., 2018). This sort of contradiction between priorities represents one core of where bias should be investigated in the system; when social workers need to consider two major goals in their execution of a process, one may very well get ignored or passed over. The bias described by Featherston was most present when social workers were searching for which set of information to base their understanding of the case on, further emphasizing the importance of intake and treatment plan development in the generalized function.

The generalized functions of case management contain the tasks one might most typically associate with this work. They consist of getting youth into the system (intake), deciding what services they will receive (development of treatment plan), executing that treatment plan (treatment), and finally moving the youth to a different program or permanent home (transfer). These functions are of extremely high importance to final youth outcomes and, as a result, the equity of those outcomes. Intake in particular is a very data heavy step, as client history, demographics, notes from previous agencies, and other information is collected about the case. This information would all exist at the physical form level of the abstraction hierarchy, and will be included in future versions of the model. However, that information then needs to be heavily filtered through a social worker's training and past experiences, possibly leading to a biased understanding of the case.

As previously mentioned, there are substantial constraints on treatment plan development based on availability of resources and the programs that are currently offered. This highlights the importance of efficient program management to ensure that the programs needed by youth are being provided. Given the high importance and numerous constraints included in the treatment plan development, this is the area that would most benefit from the addition of a decision aid.

Decision Aid Insight and Limitations

A decision aid designed to assist in treatment plan development would need to consider all of that task's dependencies. It would need to consider the running of group activities and legislative requirements, both of which may limit which programs the youth could be offered. Similarly, equitable treatment between cases and family preferences for the case should be included in the logic of the program. As all these more abstract goals are included in the development of the decision aid, more work is needed to conclude how the aid would meet those goals. The final two traditional levels of an abstraction hierarchy, physical function and physical form, will provide insight into what concrete documents, data points, and case factors drive these decisions. This data (and manipulations of it) will be what is displayed by a decision aid, and as such it is crucial to understand its place in the hierarchy.

While currently a large amount of data is collected, it is often unclear how, or even if, much of it is used for case decisions. Particularly given the heuristic decision making seen in previous studies (Zeijlmans, López López, Grietens, & Knorth, 2019), it seems that much of this data is currently underutilized. Further development of this model by way of naturalistic study of social worker data use would provide powerful insight into how this data might best be presented by a decision aid. This work could be conducted by way of work observation, subject matter expert interviews, and retroactive analysis of case data and other workplace documents.

CONCLUSION

As advancements in automation continue and more powerful decision aids become technologically feasible, it becomes more and more important to develop these aids in a manner that is contextually aware. In the foster care domain, one of the more important contexts is the presence of systemic and individual bias. In order to frame the development of a future decision aid, a formative model of the work conducted in the domain is highly necessary. This paper presents the beginnings of one such model in the form of a partial abstraction hierarchy, which captures the dependencies between the goals and constraints of the work at an ecological level.

This model was developed using insight from subject matter experts and models two major areas of work: case and program management. The program management level aims to efficiently use resources to match the agency's offerings to community needs. On the other hand, the case management level focuses on how individual youth's treatment plans are developed and executed to promote their wellbeing and meet their family's desires. There are a number of situations in which these two levels interact including limitations placed on youth treatment by program constraints and a need for extensive case documentation to allow for program funding. These areas of friction between goals are one particular source of concern for where bias may enter into the system should case workers ignore one goal for another. Additionally, social workers' information search during the intake and treatment plan development processes drive much of the youth outcome and are particularly vulnerable to individual and systemic biases.

Future decision aids should be designed in such a manner to support not just the goals of a specific task, but its

dependencies as well. This project will continue with further development of this abstraction hierarchy to include physical function and physical form levels. These levels will include the specific data points which could be used in a decision aid, and as such this development will be necessary to create a contextually-informed decision aid. While there is still work to be done, this early model has already shown a number of insights into future work and the nature of bias in the domain.

REFERENCES

116th Congress Bill Profile H.R. 2702, Family First Transition and Support Act of 2019, Bill Profile. (2019). In.

Biehal, N., Sinclair, I., & Wade, J. (2015). Reunifying abused or neglected children: Decision-making and outcomes. Child Abuse & Neglect, 49, 107-118. doi:10.1016/j.chiabu.2015.04.014

Boyas, C. T. a. M. (2017). Predictive Analytics in Child Welfare: An Assessment of Current Efforts, Challenges, and Opportunities. Retrieved from

Carden, T., Goode, N., & Salmon, P. M. (2019). Accounting for memes in sociotechnical systems: extending the abstraction hierarchy to consider cognitive objects. Ergonomics, 62(7), 849-863. doi:10.1080/00140139.2019.1603403

Chouldechova, A., Benavides-Prado, D., Fialko, O., & Vaithianathan, R. (2018). A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. Paper presented at the Proceedings of the 1st Conference on Fairness, Accountability and Transparency, Proceedings of Machine Learning Research. http://proceedings.mlr.press

Coury, B. G., & Schulze, D. (2004). Using Fault Trees and Cognitive Work Analysis to Assess Human Actions in Aviation Accidents. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 48(16), 1978-1982. doi:10.1177/154193120404801639

Dettlaff, A. J., Christopher Graham, J., Holzman, J., Baumann, D. J., & Fluke, J. D. (2015). Development of an instrument to understand the child protective services decision-making process, with a focus on placement decisions. Child Abuse & Neglect, 49, 24-34. doi:10.1016/j.chiabu.2015.04.007

Featherston, R., Shlonsky, A., Lewis, C., Luong, M.-L., Downie, L., Vogel, A., . . . Galvin, K. (2018). Interventions to Mitigate Bias in Social Work Decision-Making: A Systematic Review. *Research on Social Work Practice, 29*, 104973151881916. doi:10.1177/1049731518819160

Jiancaro, T., Jamieson, G. A., & Mihailidis, A. (2014). Twenty Years of Cognitive Work Analysis in Health Care: A Scoping Review. Journal of cognitive engineering and decision making, 8(1), 3-22. doi:10.1177/1555343413488391

McIlroy, R. C., & Stanton, N. A. (2011). Getting past first base: Going all the way with Cognitive Work Analysis. Applied ergonomics, 42(2), 358-370. doi:10.1016/j.apergo.2010.08.006

Munro, E. (1999). Common errors of reasoning in child protection work. Child Abuse Negl, 23(8), 745-758. doi:10.1016/s0145-2134(99)00053-8

Paudyal, P., & William Wong, B. L. (2018). Algorithmic Opacity: Making Algorithmic Processes Transparent through Abstraction Hierarchy. In (Vol. 62, pp. 192-196).

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). Cognitive systems engineering / Jens Rasmussen, Annelise Mark Pejtersen, L.P. Goodstein. New York: Wiley.

Read, G. J. M., Salmon, P. M., & Lenné, M. G. (2012). From work analysis to work design: A review of cognitive work analysis design applications. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 56(1), 368-372. doi:10.1177/1071181312561084

Read, G. J. M., Salmon, P. M., Lenné, M. G., & Stanton, N. A. (2015). Designing sociotechnical systems with cognitive work analysis: putting theory back into practice. Ergonomics, 58(5), 822-851. doi:10.1080/00140139.2014.980335

Rouse, W. B., Pennock, M. J., Oghbaie, M., & Liu, C. (2017). Interactive visualizations for decision support: Application of Rasmussen's abstraction-aggregation hierarchy. Applied ergonomics, 59(Pt B), 541-553. doi:10.1016/j.apergo.2016.03.006

Stanton, N. A., Salmon, P. M., Walker, G. H., Green, D., Ashleigh, M., & Baber, C. (2004). The role of the abstraction hierarchy in team performance. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 48(20), 2421-2425. doi:10.1177/154193120404802010

Suresh, H., & Guttag, J. V. (2019). A Framework for Understanding Unintended Consequences of Machine Learning.

Vicente, K. J. (1999). Cognitive work analysis toward safe, productive, and healthy computer-based work / Kim J. Vicente. Mahwah, N.J: Lawrence Erlbaum Associates.

Vicente, K. J. (2002). Ecological Interface Design: Progress and Challenges. Human factors, 44(1), 62-78. doi:10.1518/0018720024494829

Wong, W. B. L., Sallis, P. J., & O'Hare, D. (1998). The ecological approach to interface design: applying the abstraction hierarchy to intentional domains. In (pp. 144-151): IEEE.

Zeijlmans, K., López López, M., Grietens, H., & Knorth, E. J. (2019). Heuristic decision-making in foster care matching: Evidence from a think-aloud study. Child abuse & neglect, 88, 400-411. doi:10.1016/j.chiabu.2018.12.007