Children and Youth Services Review 58 (2015) 127–136

Contents lists available at ScienceDirect

Children and Youth Services Review

journal homepage: www.elsevier.com/locate/childyouth

Using big data for evidence based governance in child welfare

Hye-Chung Kum [a,c,*], C. Joy Stewart [b], Roderick A. Rose [b], Dean F. Duncan [b]

[a] School of Public Health, Texas A&M Health Science Center, 1266 TAMU, College Station, TX 77843, United States
[b] School of Social Work, University of North Carolina at Chapel Hill, 325 Pittsboro St, Chapel Hill, NC 27599-3559, United States
[c] Department of Computer Science and Engineering, Texas A&M University, 3112 TAMU, College Station, TX 77843, United States

ARTICLE INFO

Article history:
Received 1 February 2015
Received in revised form 17 September 2015
Accepted 17 September 2015
Available online 21 September 2015

Keywords:
Big data
Evidence based governance
Knowledge discovery and data mining (KDD)
Data science
Population informatics
Policy informatics
Academic government partnership
Administrative data

ABSTRACT

Numerous approaches are available for improving governance of the child welfare system, all of which require longitudinal data reporting on child welfare clients. A substantial amount of agency administrative information – big data – can be transformed into knowledge for policy and management actions through a rigorous information generation process. Important properties of the information generation process are that it must generate accurate, timely information while protecting the confidentiality of the clients. In addition, it must be extensible to serve an ever-changing policy and technology environment. Knowledge discovery and data mining (KDD), aka data science, is a method developed in the private sector to mine consumer data and can be used in public settings to support evidence based governance. KDD consists of a rigorous 5-step process that includes a Web-based end-user interface. The relationship between KDD and governance is a continuous feedback cycle that enables ongoing development of new information and knowledge as stakeholders identify emerging needs. In this paper, we synthesis the different frameworks for utilizing big data for public governance, introduce the KDD process, describe the nature of big data in child welfare, and then present an updated KDD architecture that can support these frameworks to utilize big data for governance. We also demonstrate the role KDD plays in child welfare management through 2 case studies. We conclude with a discussion on implications for agency–university partnerships and research-to-practice.

1. Introduction

An important goal of policy and management research in child welfare is to improve the governance of the child welfare system. As a term applied to public policy, *governance* refers to the act of governing (Emerson, Nabatchi, & Balogh, 2011), the social coordination of public services (Lowndes & Skelchner, 1998), or the "creation, execution, and implementation of activities backed by the shared goals of citizens" (Bingham, Nabatchi, & O'Leary, 2005, p. 548). A number of related approaches such as evidence-based management (Kovner, Elton, & Billings, 2000) or performance management (Heinrich, 2007) have the potential to improve child welfare system governance through the use of performance information to inform the design and selection of interventions by managers and caseworkers. These approaches all require measures of the experiences of child welfare clients over time and across a spectrum of desirable social outcomes. The child welfare system, like the social welfare system in broader perspective, benefits from a wealth of administrative data collected by agencies for the discharge of their duties to clients and the public.

Converting these data into performance information that can be used by agencies and other stakeholders to improve the lives of clients and the effectiveness and efficiency of the child welfare system is not a straightforward task. This task involves creating consistent measures (Usher, Locklin, Wildfire, & Harris, 2001) while serving the diverse information needs of various stakeholders, maintaining accuracy in an ever-changing policy and technology environment, and ensuring client confidentiality. Knowledge discovery and data mining (KDD) is a rigorous approach to the integration, processing, modeling and distribution of such data. In the KDD process, administrative data are transformed first into longitudinal data consisting of key social indicators at different units (e.g. county, judicial district) and then subsequently modeled into "data cubes" for retrieval through a dynamic on-demand web interface. Data cubes are multi-dimensional tables that store the various pre-calculated measures from all aspects to support easy use. In addition, the individual level longitudinal data are used to support in depth analysis for reports, research, and evaluation. The relationship between KDD and governance is not one way, but rather exists in a feedback cycle that enables the continued development of new data infrastructure and reporting tools as needs arise or are recognized by stakeholders.

In this article, we extend prior works by situating KDD as a concept important to public governance, discussing several contemporary approaches to management that public agencies can use to effectively serve their populations, introduce the basics of big data and the KDD

* Corresponding author at: School of Public Health, Texas A&M Health Science Center, 1266 TAMU, College Station, TX 77843, United States.
E-mail addresses: kum@tamhsc.edu (H.-C. Kum), jstewart@unc.edu (C. Joy Stewart), rarose@email.unc.edu (R.A. Rose), dfduncan@email.unc.edu (D.F. Duncan).

http://dx.doi.org/10.1016/j.childyouth.2015.09.014
0190-7409/

process, and then further develop the KDD architecture. Data and information technology needs are presented as central to these approaches or any approaches that purport to translate knowledge about the population served into action and monitor performance such as implementation sciences (Fixsen, Blase, Naoom, & Wallace, 2009). We describe the generic KDD process for processing big data effectively and then present our updated system architecture developed for public governance applications detailing the new elements including examples of more advanced work in analytics. We demonstrate its efficacy by describing two demonstrations of KDD system in use. We conclude with a discussion of the strong partnerships between government agencies and interdisciplinary teams at public universities that can lead to successful implementation of comprehensive KDD information systems for child welfare.

2. Public governance

Generally, governance involves the management of activities and services informed by and contributing to the public interest with such management typically undertaken by multiple stakeholders (e.g., Emerson et al., 2011; Hill & Lynn, 2004). These stakeholders can include practitioners from government agencies (e.g., public managers), non-profit organization and private for-profit firms (Salamon, 2001). In social services, the stakeholders typically include county and state departments or divisions of social service managers and caseworkers, and private non-profit organizations.

Evidence and theory suggest that governance characterized by horizontally organized and collaborative partnerships between stakeholders, such as between a government agency and a private organization, may be more effective in achieving the conditions necessary to promote socially desirable outcomes than top-down hierarchical relationships between government agencies and their subordinates (Emerson et al., 2011). Reorganizing public services into partnerships between public and private organizations, and a change in perspective from counting inputs or adhering to rules to a performance or accountability orientation are therefore central parts of public service effectiveness efforts (Heinrich, 2002). These new structures and performance orientations also have the potential to improve efficiency, offering these services at lower overall cost than their hierarchical predecessors (Benish, 2010; Head, 2008). Public policy scholars have long recognized that the effectiveness and efficiency of governance can be improved using social and behavioral science methods (Hill & Lynn, 2004). Three inter- related threads in the policy and management literature, evidence-based policy (EBP), performance management (PM), and evidence-based management (EBM), organize effectiveness efforts into cohesive frameworks that support public governance. We compare the three frameworks in Table 1, and demonstrate its relationship to KDD in Fig. 1.

3. Role of information systems in public governance

Scholars across fields such as policy (e.g., Bardach, 2003) and medicine (e.g., Sim et al., 2001) recognize the need for information systems that support the use of evidence-based best practices. Because child welfare agency managers (using, for example, PM or EBM) must make plans on the basis of performance and have ways to monitor the

Table 1
Comparison of public governance frameworks.

	Evidence-based policy (EBP)	Performance management (PM)	Evidence-based management (EBM)
Description	Policy decisions implemented on the basis of scientific research & evaluation.	Practices, technology, and information agencies use to monitor performance; decisions informed by monitoring in the context of goals; and actions taken to achieve goals[5,8]	"The conscientious, explicit, and judicious use of current best reasoning and experience in making decisions about strategic interventions"[11]
Goal	Identify and institutional effective policy interventions through the use of experiments[1]	Generate timely institutional knowledge through the day-to-day activities of governance[8]	Improve organizational performance through trial-and-error experimentation and continuous evaluation of staff activities and program outcomes[12]
Practice	Agencies can use extant literature to choose the most effective and efficient practices for addressing the unique challenges identified by data on agency performance[2]	Agencies make the most informed decisions possible that support their objectives, even in the absence of clear scientific evidence[9]; agencies conduct performance monitoring[8]	Multi-step processes: Scientific management, continuous quality improvement, and organizational learning; adapt, combine, and discontinue programs in response to changing agency conditions and client information[12]
Requirements	Information that identifies areas of need; access to evidence connecting these needs to appropriate policy interventions	Timely and accurate information technology structures[10]; organizational capacity to interpret data	Accumulation of long term data trends and information technology structures that handle these data[10]; organizational capacity and partnerships with research organizations to interpret data and change structures accordingly[10]
Challenges	Evidence and theory cannot keep pace with the needs of practitioners; agencies must perform certain tasks without evidence of their effectiveness[3,4]; Practices are partly informed by values of practitioners[1,5]; goals change over time[6,7]	Evidence may not be scientifically valid without controls or comparisons to a hypothesized counterfactual; agencies may not be staffed to handle the data visualization and analysis needs	Evidence may not be scientifically valid without controls or comparisons to a hypothesized counterfactual; agencies may not be staffed or have the partnerships needed to handle the data visualization and analysis needs
Relationship to Other Approaches	None inherent	Complementary to evidence-based policy[8]	The methods and processes of PM subjected to scientific scrutiny[13]
Role of KDD	Identify needs	Identify needs and monitor and evaluate program implementation and outcomes	Identify needs, monitor and evaluate program implementation and outcomes

[1] Jennings & Hall, 2004.
[2] Cannon & Kilburn, 2003.
[3] Bingham et al., 2005.
[4] Bardach, 2003.
[5] Head, 2008.
[6] Benish, 2010.
[7] Maynard, 2006.
[8] Heinrich, 2007.
[9] Schorr & Auspos, 2003.
[10] Meier & O'Toole, 2009.
[11] Kovner et al., 2000, p. 10.
[12] McBeath, Briggs, & Aisenberg, 2009.
[13] Briggs & McBeath, 2009.

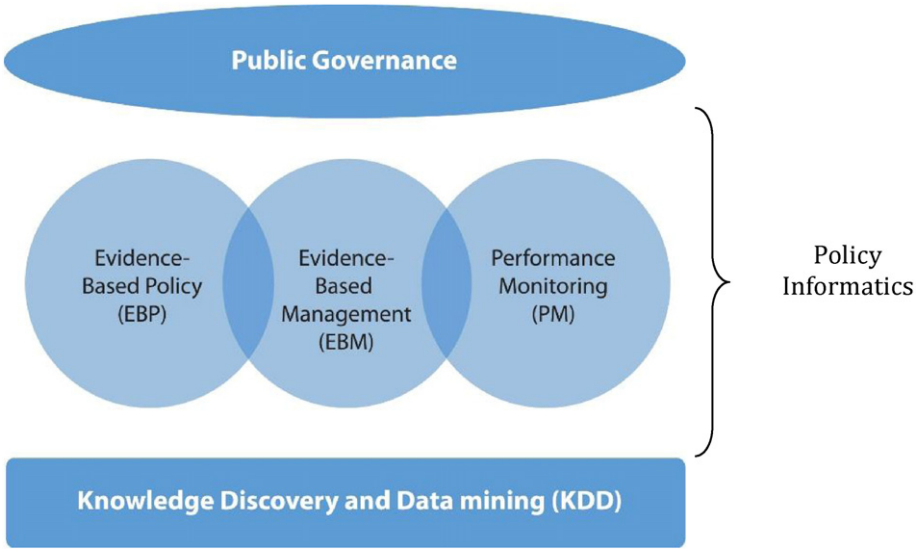

Fig. 1. KDD supports frameworks to increase effectiveness of public governance.

performance of their actions, they must also have access to information systems with the capacity to report on client and system performance (i.e., child welfare informatics; Naccarato, 2010; Nguyen, 2007). More generally, efforts to improve governance—whether through decentralization of authority to private partners or strictly hierarchical forms—require some type of feedback-enabling system in which the status quo can be measured, actions taken, and performance associated with those actions observed (Ammons & Rivenbark, 2008; Berman & Wang, 2000; Bingham et al., 2005). These requirements imply specific types of data. Descriptive data including outcomes that indicate system performance and individual well-being are needed for targeted planning, contract development, and to assess performance during a contract period (e.g., with private firms; Flaherty, Collins-Camargo, & Lee, 2008). Further, tracking agency performance over time, such as through event histories, requires measures taken over time using panel or longitudinal designs consisting of multiple measures on the same family or client (Poertner, Moore, & McDonald, 2008).

Although strong evidence supports the importance of these information systems to effective governance, there may be problems with realizing such systems. Issues with the capacity for performance-related reform were raised by Berman and Wang (2000). More recent qualitative research suggests that a typical county social service agency may still experience certain barriers, such as not having the capacity to monitor the status of or report on their children, and that the coordination necessary for the construction of useful data may be burdensome to agency staff (Meezan & McBeath, 2011). For example, agencies typically collect their own administrative data for the purpose of conducting their operations but lack the means to prepare the data in ways that are meaningful to governance, such as combining the data with interoperable sources from multiple arms of the same agency or other agencies in the same county, or tracking clients over time (Penuel & Means, 2011). Further, without forming cross-county partnerships themselves, county agencies lack access to similar data in other counties that they can use as performance benchmarks (e.g., by comparing themselves with counties of similar sizes or populations). These barriers suggest the potential benefits for a more centralized service that local agencies can utilize, such as a state-level repository of prepared data, eliminating the need for burdensome inter-county partnerships. A relationship between county and state social service agencies complemented by a centralized informatics system can support and promote effective governance if the data can be utilized by county agencies in ways that promote the well-being of clients. In addition, the need for high level expertise and cost of developing a system for one county is almost as much as doing so for all counties.

Further, issues related to security, consistency, diversity, and accuracy of the data must be addressed by these methods. First, allowing multiple counties access to statewide raw administrative data raises confidentiality concerns as individual clients can be re-identified through such data (e.g., staff from one county could identify clients in a neighboring county). Descriptive modeling of the raw data and archiving of these data in secure environment (Kum & Ahalt, 2013, March) prevent counties from reporting data that re-identifies individual clients. Second, the outcomes reported by agencies should be consistent; that is, each measure means the same thing to all stakeholders, representing important social metrics of progress in child welfare management. Third, diversity implies that a number of different measures or ways of looking at the data are possible, and that stakeholders' unique needs are represented. Finally, the data should be clean and as accurate as possible. Typical approaches to protecting confidentiality such as top-coding of individual-level data (i.e., lumping the tails of the distribution) may render the data inaccurate, potentially biasing descriptive statistics (Lane, 2005). KDD is an effective means of constructing and distributing data in support of the frameworks to increase the effectiveness of public governance (see Fig. 1). KDD addresses the limitations noted above by helping individual agencies bypass processing data themselves and gives them access to quality data on their performance as well as that of their peers.

4. Big data, knowledge discovery and data mining (KDD), and data science

Informatics is the science of information, the practice of information processing, and the engineering of information systems. Policy informatics is an approach to research and practice of managing complex policy issues through these systems. It is a subfield of the emerging field of population informatics, the study of populations using big data. Informatics studies focus on the technological methods and user behaviors, as well as the interactions between technology and users in the manipulation and processing of quality information for addressing policy and managerial problems. Manually processing data on a one-time basis, common in the social sciences, is not sufficient to address the information required for policy informatics. Rather, ongoing processes are a better fit for policy informatics. In this section, we present how to develop an efficient information processing pipeline to support policy informatics.

4.1. What is knowledge discovery and data mining (KDD) and data science?

KDD is a framework that can be used to engineer such an information system. KDD refers to the process of finding knowledge in data to promote timely action. KDD was initially developed for private companies to better manage and market through predictive modeling (Kum, Duncan, Flair, & Wang, 2003). However, KDD is a rigorous and systematic method for processing data in a flexible and tractable manner beyond predictive modeling, and provides an appropriate framework for processing data for policy informatics (Hand, Mannila, & Smyth, 2001).

KDD combines ideas from multiple traditions—databases, machine learning, artificial intelligence, knowledge-based systems, information retrieval, statistics, pattern recognition, and visualization—yielding an integrated approach to extracting valuable information from data. KDD is "the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data," (Fayyad, Piatetsky-Shapiro, & Smyth, 1996, p.30). It enables the conversion and interpretation of multifaceted raw data into knowledge that can be easily consumed by users (Fayyad et al., 1996c). In addition, an important aspect of KDD is the timely delivery of this information to those who will ultimately take action, which is important to effective governance models. KDD technology allows for both consistency and diversity throughout the data. To effectively address the many local problems faced by child welfare agencies, diversity and consistency are required (Duncan, Kum, Weigensberg, Flair, & Stewart, 2008).

Fig. 2 depicts the typical elements of the generic KDD process. The first step is to ingest the operational data into an Enterprise Data Warehouse (EDW) through an ETL (Extract-Transform-Load) process which involves cleaning, integrating, and developing new calculated variables that can measure important concepts. Often these new measures are called features in the computer science and statistics literature. Typically EDW are in a SQL database, but that is not necessary. The EDW are similar to the base longitudinal files often developed in the child welfare field. Second, from the EDW the required variables for a certain problem are selected for the task specific dataset for further processing. These datasets are often called the analytic dataset in child welfare, and represent the table that is fed into a typical regression model. Third, standard statistics, advanced analytics and data mining algorithms are applied to this analytic dataset to generate results to address a specific problem. This step is the most published step in the KDD process, and many people confuse this step with the full KDD process. But we note that modeling is only one step in the full process of converting data into actionable information. One common reason for failed projects is too much focus on modeling and a lack of attention to the rest of the process. Often many projects fail due to insufficient relevant features, lack of appropriate data among what seems abundance of data, or insufficient expertise in implementation rather than lack of modeling experts. Having expertise in the full process is key to successfully using data for governance. Fourth, the results are evaluated, often taking the form of validation against a test dataset that was held out during the modeling phase. If the results are not satisfactory, you iterate back to one of the earlier step, to adjust the process where appropriate. Processing big data into actionable information is a very iterative process, where the analysis is refined over time as the researcher interacts with and learns the data. Finally, when the results are acceptable, they are converted into the appropriate information form (e.g., papers, reports, presentations, website) and delivered in a timely manner to facilitate decision and action.

The various computer programs that convert the raw data to information need to be designed as an incremental flow system that can be maintained with little effort. The theme of designing information systems as processes, which can handle the constant flow of data, is key to understanding how to effectively use big data for governance. These longitudinal data files can be analyzed for various reports, research, and monitoring of performance measures, which is more comprehensive, accurate, and easier to use than individual raw tables. KDD systems must be agile and extensible, so that they can adapt to changes in programs, environments, and targets over time. As a case in point, such changes represent one of the major challenges that limit the application of EBP, suggesting PM as an alternative and complementary approach. KDD systems should be adaptable to monitor changes in (for example) programs and produce current information in a timely

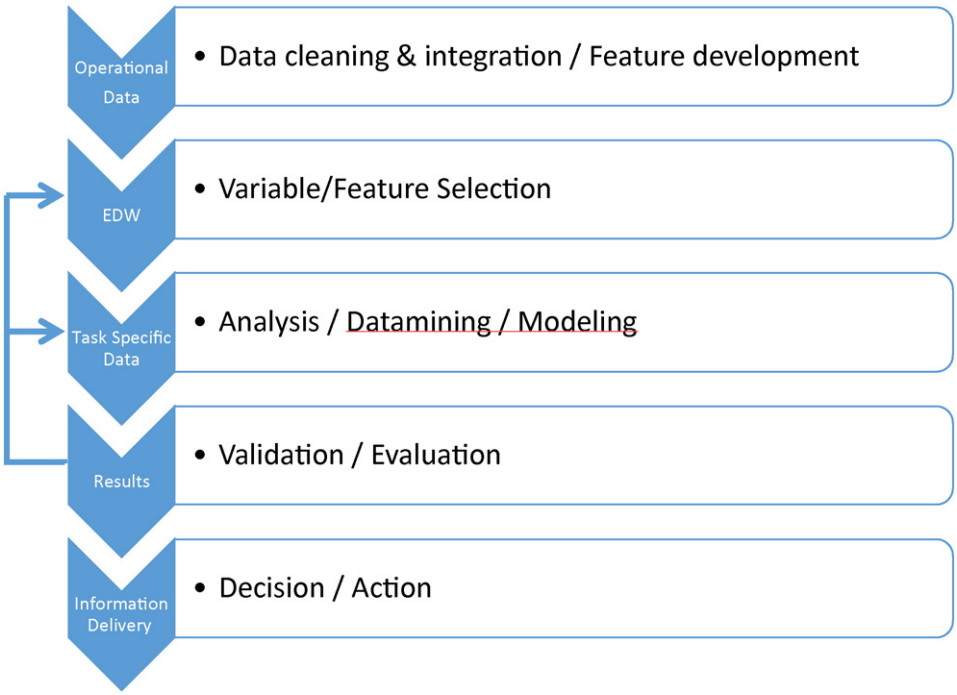

Fig. 2. KDD Process.

manner (e.g., compare performance under new programs with performance under older programs); provide useful program performance information in an unstable environment (e.g., fluctuations in the economy; changes in other agencies); and adapt to changing performance targets (e.g., a capacity for adding new data). Extensible system design allows for the agility required in such an information system. Thus, KDD information systems support processes, providing more immediate and comprehensive information than is typical in one-time evaluations.

We note that recently, with the significant interest in big data, loosely the term "data science" has become much more popular among the general public to refer to the KDD process which was mainly used among experts in the field. In this paper, we use the more technical and clearly defined term KDD.

4.2. What is big data?

A central characteristic of KDD is that it translates operational data that are primarily collected for the purpose of managing an organization into performance information on individual clients, families and populations. In child welfare, these operational data are administrative data collected by child welfare and other government agencies, and include findings of investigations of child abuse and neglect, foster care placement types (such as licensed foster care or kinship homes), and length of stay in foster care in the population (rather than a sample). In sum, important characteristics of administrative data are that (1) although fragmented, an extensive amount of data exists on the full population served, (2) they are continuously generated, (3) they change over time as programs evolve and originate from multiple sources, and (4) they have varied levels of validity with data directly required for operation being the most valid. These represent the "four Vs" of big data: volume, velocity, variety, and veracity, respectively.

Some experts in big data also discuss a fifth V of big data as being "value". In the context of using administrative data for governance, it can be interpreted as the potential value of administrative data (for example, for PM or EBM), when properly processed into information and knowledge to support good decisions and actions in policy and management. Hence, big data is much more than large volumes of data and in fact, in child welfare velocity, variety, and veracity pose the most challenges and these challenges makes integrating large volumes of coherent data difficult. A well-designed KDD information system on child welfare is well suited to tackle these challenges. It can greatly enhance federal, state, and local public governance as KDD can effectively process the chaotic data into useful information and facilitate the sharing of timely relevant information between many diverse stakeholders from diverse sources.

4.3. An updated KDD system architecture for public governance

Important components of a KDD information system for governance (e.g., in child welfare) include (1) methods for accurately and incrementally integrating data from different sources on a regular basis, (2) a flexible model for organizing the data to make it easy to use for longitudinal analysis but also have tolerance for error and change present in all real data, (3) an efficient mechanism to disseminate information that has been extracted, often through more than one channel, and (4) a secure architecture that provides confidentiality protection to the subjects of the data.

KDD makes possible consumer- and client-oriented information that was not possible before by generating actionable information in a timely manner. For example, the KDD system architecture is used at the supermarket to generated "personalized coupons". Processing raw data into actionable information requires a systematic continuous processing of big data that is only possible using a framework such as KDD. In the context of public governance, KDD is a systematic framework for processing big data that provides timely information for action in policy and management. In the next section, we describe the details of a working KDD system administered by a university for serving county child welfare agencies and provide several examples of the use of this system.

4.3.1. The full KDD system architecture

The primary objective of our KDD information system is to inform welfare policy and practice in North Carolina (Kum et al., 2003). It was originally funded in 1997 through a contract with the NC Division of Social Services (NC-DSS) to evaluate WorkFirst (North Carolina's Temporary Assistance for Needy Families program) under welfare reform. The system has since grown to cover Work First, Food and Nutrition Services (formerly known as Food Stamps), and Child Welfare. In this paper, we limit our discussion to child welfare. Child welfare agencies typically have administrative records pertaining to the experiences of children and families that were served by the child welfare system. These data can be used to consistently and accurately depict the diverse circumstances of children in the child welfare system.

The basic system for hosting a website was detailed in Kum, Duncan, & Stewart, 2009. In this paper, we update the system to include data from other sources as well as use of the data infrastructure for other purposes. Fig. 3 illustrates the full KDD system built for the NC project. This figure shows the full cycle of the KDD process, from data acquisition to consumption of the information via a public website as well as ad hoc reports, presentations, and papers. First, the data regarding child welfare cases are acquired as raw administrative data tables from the state division of social services (NC-DSS).

We note that these data are entered into a state-operated database system by county social workers from all 100 counties in North Carolina as the program is administered. Second, the development team weaves these raw tables together to build longitudinal files, depicting the experiences of each child over time as child welfare measures. Relevant data from other systems that widen the scope or add the appropriate context (i.e. educational experiences of children or employment outcomes of youth aging out of care) will also be added during this step. Depending on the relationship with other agencies, outside data are fed into the KDD system continuously or on a one-time basis. The computer programs that weave the files together need to support monthly incremental additions over time with little effort. Third, descriptive models of these measures are stored in summary tables, called data cubes, for easy access. Fourth, these data cubes are shared with all 100 counties and other stakeholders of the child welfare system via a public website that enables users to examine the data in a variety of dimensions. The dynamic website can generate county reports and graphs based on users' selections from a series of menus. From the website, local agency staff can easily obtain timely measures regarding their performance by monitoring changes in children's outcomes over time as well as across counties. Fifth, information delivery to agency staff is facilitated by documentation, training and responses to data requests. These five steps are supplemented and supported by a rigorous software engineering process.

4.3.2. Descriptive modeling

In this section, we summarize the key concepts for the first stage in developing the data infrastructure for public governance. Ideally, the KDD process works best when data are generated and consumed by the same users, which in this case would be county child welfare agency staff. This full cycle of data generation and consumption ensures the quality of data entered into the system as well as proper interpretation of the data elements. Thus, the first stage of implementing the data infrastructure is developing a descriptive modeling system that is useful for those that are administering the child welfare programs and generating the data. A common system involves a dynamic website where descriptive statistics are published about the data. The process of developing the useful measures to publish on such a website, and having users who generate the data start to use information produced from their own data ensures that the data is either valid for use or will

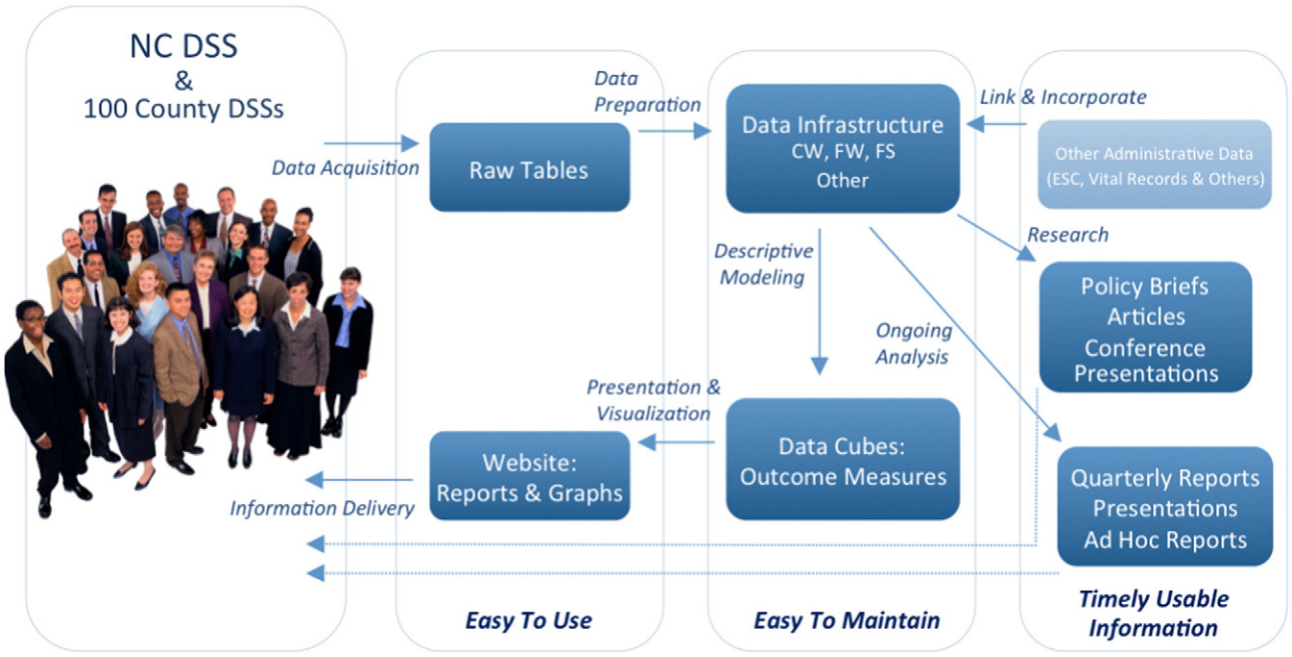

Fig. 3. KDD information system architecture.

become valid for use quickly because the data producers will correct the problems in the data to obtain the information needed as they consume the information. Thus, completing the data cycle such that the data flows back to those that enter the data is the first stage. For details on implementation see Kum et al., 2009.

The purpose of descriptive modeling in KDD is to extract compact and easily understood information from large data files (Hand et al., 2001). Thus, in the context of governance the goal of this step is to use the data warehouse to calculate measures that are meaningful to stakeholders and that inform performance. The main products of descriptive modeling are data cubes that are multi-dimensional tables that store the various pre-calculated measures from all aspects (e.g., different units, different time periods). Data cubes can be quickly called by the web-based user interface that enable examination of the longitudinal data from multiple dimensions and perspectives. In the NC KDD project, the main goal was descriptive modeling to inform child welfare policy and practice in North Carolina. The NC KDD project uses a collection of outcome measures to accurately describe the experiences of children in child welfare in each county and statewide (Duncan et al., 2008). Taking into account the history of outcome measures in child welfare, a combination of longitudinal and cross sectional measures on placement in foster care as well as reports of abuse or neglect were used. A detailed list of measures used and sample screenshots of the website can be found in Duncan et al., 2008.

Several factors need to be considered in descriptive modeling, including choosing the unit of analysis, selection of the measure and method used to define and calculate the outcome, and the period of time covered by the measure. In the context of governance and PM, these factors assume that an outcome or set of outcomes has been identified through either a careful review of the agency's mission statement, by other clearly defined program objectives, by statutory requirements, or by decisions made in the political system such as through the interactions of managers with caseworkers or legislators. In sum, via the dynamic website each county child welfare agency can view their performance from various perspectives using a set of outcomes depicted from different unit, measure, and time perspectives. Subgroup analyses such as for age, gender, race, and ethnicity, can also be reported. The design ensures that agencies have a comprehensive view of the child welfare system in North Carolina.

4.3.3. Beyond descriptive modeling.

The web-based delivery system is not the only information channel through which the data are shared with public managers. Once the data have been curated into comprehensive longitudinal analysis files, they can be used in two other ways that may ultimately be fed back into the descriptive modeling system as new data cubes if the need or opportunity arises. The first use consists of ongoing and ad hoc analysis performed on behalf of our institutional partners at all government levels. These ad-hoc reports have much more information than what can be cleanly defined in the data cubes described above. In addition, we have generated quarterly reports that assess changes in the monthly child welfare caseload over time and provide in-depth analysis of special topics in child welfare, such as variation in exit type by age (Duncan, 2010) and county-level trends in the number and rate of children in foster care (Stewart & Duncan, 2013). One-time evaluations and studies are possible and have been conducted using this system. Even these studies are easily updated with more current data as it becomes available when useful.

The second is the use of these files for writing papers, proposals, research, and evaluation. The data have been used in several doctoral dissertations and we have assisted a few researchers outside of our team with using the data files. A few examples are noted below;

- Lawrence, Rosanbalm, and Dodge (2011) used the longitudinal dataset of child abuse and neglect reports in an evaluation of North Carolina's alternative response system to reports of abuse and neglect.
- Stewart, Kum, Barth, and Duncan (2014) updated an earlier technical report from 2008 to track employment outcomes up to age 30 for youth who aged out of foster care. This work would not have been possible if not for the KDD system that made the updated analysis relatively simple.
- We have used the KDD system for economic forecasting. In (Duncan & Stewart, 2007), caseload information was used to project child welfare and economic services staffing needs for Meckenburg County. In (Duncan, 2007), data on child welfare caseloads and expenditures was used to project changes in cost neutrality that led to North

Carolina ending its IV-E waiver initiative early.

- The KDD system has also been used to support advanced algorithm development in sequential analysis (Kum, Duncan, & Wang, 2004) and privacy preserving interactive record linkage (Kum, Krishnamurthy, Machanavajjhala, Reiter, & Ahalt, 2014). In the sequential analysis, we developed an algorithm to detect patterns in sequences of sets to find common patterns of service use over time. This allowed us to investigate monthly patterns of service use (a set of services) over time for children with substantiated reports of abuse and neglect. In the most recent algorithmic work in privacy preserving interactive record linkage, we designed a computer-based third-party linkage platform that can precisely control the information disclosed during the record linkage phase to significantly improve the quality of data integration while reducing the privacy risk to negligible levels when linking sensitive data.
- Most recently, we are applying propensity score matching (PSM) to foster care and Medicaid data to compare outcomes at the county level in the evaluation of a US Children's Bureau grant exploring moving toward a trauma-informed child welfare system. PSM is often used to balance measured covariates to find comparison matching groups in observation studies. However, if the unmeasured covariates affecting the intervention are confounders, PSM can exacerbate the bias in estimates, and thus must be used with caution (Brooks & Ohsfeldt, 2013).

All papers and proposals based on these data are shared with NC-DSS before final publication to give them a chance to respond and provide feedback.

4.3.4. Information delivery

Processing raw data into useful information and then presenting it is not enough to translate data into action because people have varied skill levels with data. Thus, to truly see an impact, additional steps are required to facilitate the use of the information. As a final step, the development team provides documentation and training for the website, and fills data requests from and regularly interacts with partners in state and county agencies in North Carolina. First, documentation on the details of how the data elements, entered by county staff, were processed into measures on the website is important. Because of the ongoing changes to the data and website, maintaining accurate documentation is a challenge. A page with FAQs (frequently asked questions) offered on the web site is periodically reviewed and edited for clarity and accuracy as the system changes.

Second, periodic training sessions are held with county social service agency supervisors and managers, demonstrating how to use data to track outcomes for children involved in the child welfare system and how to use the data to assess county performance on an array of indicators. In addition, the project team conducts workshops and training sessions at statewide meetings for human service professionals.

Third, over the 15 years since its inception, the website has attracted regular and advanced users who understand the measures, which provides for important two-way interactions with the development team. The development team helps users who have detailed data questions as well as request for adding more information to the site. The university development team frequently assists policy analysts and program managers from NC-DSS on various analyses of the administrative data using the data warehouse built in the KDD system. Additional analyses requests come from counties, federal agencies, legislators, and independent researchers. The data warehouse is regularly used by researchers and evaluators at the university, and represents an important asset in attracting federal, state and private funding for research and evaluation on programs and policies affecting the populations described by the data warehouse.

In addition, the interactions with users of the system provide foci for the further development and refinement of the data warehouse and data cubes. Users provide ideas that lead to new derived variables, new data cubes, and new ways of presenting the data on the website. For example, the website includes the number of children by type of child protective services assessment finding (such as, abuse, neglect, dependency, services needed and so on) organized by the state fiscal year. In the point-in-time data, children may have multiple findings associated with one assessment. We received feedback from a number of users that they wished to sum children by finding types, in particular adding abuse, neglect, and dependency together to produce a "substantiated" category. However, summing these findings produced a duplicated count of children. We responded by designing a data cube with one finding per child based on the most severe finding type and, thus, an exclusive count of children by finding type.

Taking this a step further, the KDD process provides critical end-user feedback. As users access and utilize their data, they will be motivated to submit more accurate and detailed data, which in turn would result in more reliable and comprehensive information for performance measures and outcomes. In fact, in North Carolina, when the counties saw their own outcome measures as being inaccurate as a result of bad data, many took steps to fix previous data as well as ensure future data quality. For example, with data on children's initial placement when entering foster care, the missing data for this measure constituted nearly 10% of cases in state fiscal year 1998 whereas today the percentage is almost zero. Furthermore, those at the county agencies who provide the data know best how to interpret the results, given their familiarity with agency policy and practice regarding coding and classification of data elements (Duncan et al., 2008).

4.3.5. Software engineering

Building and maintaining a KDD information system over multiple years is a complicated task that requires a systematic approach. Applying a fairly standard spiral (or iterative) model for software engineering process is the most effective method (Boehm, 1988). In an academic setting, each spiral (which includes the full development cycle of requirements, design, build and test) is typically completed within a year with much of the development done over the summer to allow students to be involved in development. Developers are required to document all phases to assure that the process can continue in the absence of the original developers (Brooks, 1975). Such assumptions significantly increase the development cost, but are required in order to keep a stable system operating in a university setting. Due to the limited staffing of the project, developers also act as testers. For this reason, we purposely build in a significant time for beta and user testing. The main user group testing the website is members of the team who are social work researchers and students. Feedback from these users not only help us find programming errors, but also significantly enhance the usability of the next iteration of the code.

5. Demonstrations: KDD in case studies of effective governance

In this section we describe two demonstrations of the use of KDD in promoting the effectiveness of child welfare agencies. These demonstrations show that KDD processes can be useful to or informed by "top-down" hierarchically organized governance structures, or "bottom-up" horizontally organized structures.

5.1. Case study 1: reaching for excellence and accountability in practice

Reaching for Excellence and Accountability in Practice (REAP), a project undertaken by a workgroup comprising state and county social service agencies, the university research and development team, and several non-profits, demonstrates the value of KDD in child welfare service governance. The REAP project began in 2009 as an effort by NC-DSS to transform the division's provision of technical assistance to county departments of social services to support continuous practice improvement (NC-DSS, 2011).

The goal of REAP was to improve the child welfare system by defining child and family outcomes and the processes and supports needed to achieve those outcomes (Bland, Johnson, Raymond, Shook, & Rozeff, 2010). Using the KDD information system built and maintained by researchers at the university in collaboration with NC-DSS, the researchers were able to support the REAP efforts by providing a list of potential outcome measures at the early workgroup meetings in a timely manner. This supported the workgroup to carry an effective conversation about final measures for REAP. Once measures were defined, the KDD system once again was used to quickly develop a dashboard tailored to the needs identified by the REAP workgroup. The data cubes in the KDD information system allowed for easily using integrated administrative data to facilitate the discussion and then to deploy the required measures to 100 diverse counties quickly. The REAP Data Dashboard outcomes ("core child welfare achievements") include measures in the areas of prevention, child protective services, foster care, permanency, reunification, adoption, transition to adulthood as well as systematic factors such as staffing and training.

5.2. Case study 2: leading by results

In contrast to the top-down model of the REAP project, the Leading by Results project used more of a bottom up approach with individual counties initiating efforts in their own county. LBR began in North Carolina in 2002 as a joint effort between the NC Association of County Directors of Social Services and the NC Association of the County Boards of Social Services. The purpose of this initiative was to improve performance by focusing on outcomes for families and individuals instead of focusing on process measures such as the time required to process a case or eligibility error rates. A task force, comprised of members of both organizations as well as the state Department of Health and Human Services, was formed in 2003 to move the effort ahead. By 2007, 28 counties were involved with LBR.

LBR has six key components, (a) developing a leadership focus, (b) developing a framework of results and indicators, (c) utilizing results data to track performance on results and indicators, (d) aligning resources and efforts to move indicators, (e) engaging and collaborating with community stakeholders in the alignment process, (f) creating a culture that focuses on outcomes. Counties involved in the initiative included the state's largest metropolitan areas as well as a number of medium-size and small counties. Although formal support for LBR declined over time, many counties are actively pursuing an LBR strategy. The main role of the KDD development team was to assist counties in using the information that is already available on the website to measure the outcomes of interest that met their local need.

6. Discussion

In this paper, we describe a method for generating child welfare agency knowledge from administrative data through a rigorous process known as KDD. KDD ensures that end-users in child welfare agencies around North Carolina have access to rich information that is sufficiently secure to prevent re-identification of clients; consistent such that two counties sharing information can be assured that their respective data represent the same underlying phenomena; and diverse such that outcomes are meaningful to a wide group of stakeholders. We start by comparing the different frameworks for evidence based governance and situate the role of information systems in these frameworks. We then introduce the basic concepts of the KDD process, not to be confused with data mining, and big data, following with a presentation of an updated KDD system architecture and a brief description of some of our more advance analytics work. We conclude with two brief case studies demonstrating the role that KDD can play in governance in the child welfare system. In the following sections we describe two key implications for management.

6.1. University-agency partnership

Strong partnerships between government agencies and interdisciplinary teams at public universities can lead to successful implementation of comprehensive KDD information systems for child welfare while providing a priceless opportunity for research. Public universities are the natural homes for such systems because (a) they are under the public oversight of state legislatures who are ultimately responsible for policies that govern state agency data, (b) they have access to content experts as well as information system experts required for building and maintaining such a system, (c) they have the flexibility and scale that most non-profit organizations or government agencies do not have, (d) the potential of the data system can be maximized and leveraged by giving researchers in the content area and information systems full access, (e) it is easier to broker relationships and integrate data from different agencies in the neutral university system with proper oversight for confidentiality, and (f) they can leverage the training of the next generation of government information specialist and scientists who will be versed in real data, technology, and data science to build and maintain these systems cost effectively.

Accordingly, the partnership between the public agencies and the university is key to effectiveness efforts such as EBM. EBM demands resources, such as sophisticated performance reporting systems, and personnel with research-related competencies that public agencies may find difficult to obtain (Briggs & McBeath, 2009). The REAP project is an example of a successful university-agency partnership that supports EBM. As noted earlier, the KDD information system developed for NC-DHHS provides the performance data needed for the REAP project. In addition, university researchers assist county agencies with critically assessing performance data and developing and monitoring their achievement plans as seen in the LBR efforts. This assistance is provided through trainings, seminars, and consultations.

The university-based research and development team is comprised of individuals with diverse backgrounds in computer science, information technology, public policy, political science and social work. The team includes members with experience providing direct services to clients and working in county departments of social services. The expertise that exists within this research team would be difficult to replicate in every county social service agency. In addition, the university infrastructure provides access to the latest advances that may be difficult for individual agencies to obtain independently. Fortunately, the partnership between the university and NC-DSS provides a mechanism for local agencies to connect with university resources and personnel. Of course at the heart of all this is a trusting relationship that has been developed over many years.

6.2. Research to practice

The relationship between the university and agencies is a mutually beneficial partnership that creates a bridge between research and practice. Agency partners, including front-line workers, supervisors and managers, provide critical information about current child welfare policy and practice and contextual factors to university researchers. This information aids researchers in generating new ideas and developing research questions and hypotheses for testing. Eventually, this process leads to research studies that analyze policy, evaluate programs, and develop and test new interventions. Results of these studies can then be fed back into EBM and EBP to improve child welfare services. The existing university-agency partnership helps to facilitate this cycle.

Over the years, the research and development team has produced or assisted other researchers in producing many studies with wide-ranging relevance to the child welfare field. These studies include, among others, felonious arrest rates of former foster youth (Barth, Duncan, Hodorowicz, & Kum, 2010); employment outcomes of youth who aged out of foster care (Macomber et al., 2008); Title IV-E Waiver evaluations in North Carolina (Usher et al., 2002; Wildfire et al.,

2007); and an investigation of the effects of the recession on child maltreatment rates (Millett, Lanier, & Drake, 2011). In addition, there were many more internal reports and dissertations covering topics such as education outcomes of children in foster care, a cost model of foster care, and the amount of duplicate IDs in the child welfare data system and its impact on the reentry rate measure.

In addition to publications and reports, data from our website and KDD information system has been used directly by other researchers such as the 2011 MRS evaluation carried out by the center for Child and Family Policy at Duke University and Kids Count, a national source of information about children maintained by Annie E. Casey Foundation. In NC, Action for Children, a non-profit organization, has a grant from Annie E. Casey to provide data about NC. The main researcher responsible for this work relies on data from our website whenever possible, and we will sometimes carry out specialized requests for information that is not published on the website using the KDD information system. Moreover, local media have used data from our website to report on child welfare.

7. Conclusion

Only data that are used regularly are valid. Most government data are rarely used and quickly end up in the piles of useless legacy data. Good decision support systems for local and state agencies under a PM or EBM framework must be deployed at the back end of the government information systems to divert such data. KDD makes these decision support systems possible. From our work, we found three models of access valuable to key stakeholders. First a public dynamic website should provide comprehensive summary statistics that are of general interest. These would naturally become the open data supporting more transparent government. Open data is a priority of the Obama administration to improve access to government data, and Data.gov is an effort to improve access to federal data. Second, a private login based individual level data should be provided to approved personnel for drill down capability. Training and changes in organization culture in agencies to use data in their daily jobs will be just as important as making the data available. When data get incorporated into the daily activities of local agencies that generate the data, administrative data will become much more valid for other purposes. Finally, a secure federated multiagency data system with confidentiality protection should be available for approved use in policy analysis and research.

The key factors to successfully implementing the KDD information infrastructure in government include trust, real support through policies and funding, access to good technical expertise in both the content area and information technology, and training (Kum, Duncan, Flair, & Wang, 2004). Access to the required expertise can be obtained through strong partnerships between government agencies and interdisciplinary teams at universities. The partnerships can lead to successful implementation of KDD information systems in the public sector while providing a priceless opportunity for research.

Acknowledgment

This research was supported in part by funding from the North Carolina Department of Health and Human Services. We also thank the reviewers for their constructive suggestions that have improved the paper.

References

Ammons, D. N., & Rivenbark, W. C. (2008). Factors influencing the use of performance data to improve municipal services: Evidence from the North Carolina Benchmarking Project. *Public Administration Review, 68,* 304–318.

Bardach, E. (2003). Creating compendia of "best practice". *Journal of Policy Analysis and Management*, *22*, 661–665. http://dx.doi.org/10.1002/pam.10160.

Barth, R. P., Duncan, D. F., Hodorowicz, M. T., & Kum, H. -C. (2010). Felonious arrests of former foster care and TANF-involved youth. *Journal of the Society for Social Work and Research*, *2*, 104–123. http://dx.doi.org/10.5243/jsswr.2010.9.

Benish, A. (2010). Re-bureaucratizing welfare administration. *Social Service Review*, *84*, 77–101. http://dx.doi.org/10.1086/653454.

Berman, E., & Wang, X. (2000). Performance measurement in U.S. counties: Capacity for reform. *Public Administration Review*, *60*, 409–420.

Bingham, L. B., Nabatchi, T., & O'Leary, R. (2005). The new governance: Practices and processes for stakeholder and citizen participation in the work of government. *Public Administration Review*, *65*, 547–558.

Bland, R., Johnson, C., Raymond, J., Shook, T., & Rozeff, L. (2010). *North Carolina REAP: Reaching Excellence and Accountability in Practice.* (Retrieved from Atlantic Coast Child Welfare Implementation Center website: http://www.accwic.org/news_events/PDFs/2010_Files/North%20Carolina%20Presentation.pdf).

Boehm, B. (1988). A spiral model of software development and enhancement. *IEEE Computer*, *21*, 61–72. http://dx.doi.org/10.1109/2.59.

Briggs, H. E., & McBeath, B. (2009). Evidence-based management: Origins, challenges, and implications for social service administration. *Administration in Social Work*, *33*, 242–261. http://dx.doi.org/10.1080/03643100902987556.

Brooks, F. P., Jr. (1975). *The mythical man-month: Essays on software engineering. reading.* Boston, MA: Addison-Wesley.

Brooks, J. M., & Ohsfeldt, R. L. (2013). Squeezing the balloon: Propensity scores and unmeasured covariate balance. *Health Services Research*, *48*(4), 1487–1507. http://dx.doi.org/10.1111/1475-6773.12020.

Cannon, J. S., & Kilburn, M. R. (2003). Meeting decision makers' needs for evidence-based information on child and family policy. *Journal of Policy Analysis and Management*, *22*, 665–669. http://dx.doi.org/10.1002/pam.10161.

Duncan, D. F. (2007). *Revised waiver expenditure forecasts.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina School of Social Work.

Duncan, D. F. (2010). *Changes and trends in the child welfare caseload in North Carolina: March 2010.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina School of Social Work.

Duncan, D. F., & Stewart, C. J. (2007). *Mecklenburg county department of social services caseload and workforce forecast.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina School of Social Work.

Duncan, D. F., Kum, H. -C., Weigensberg, E. C., Flair, K. A., & Stewart, C. J. (2008). Informing child welfare policy and practice: Using knowledge discovery and data mining technology via a dynamic web site. *Child Maltreatment*, *13*, 383–391.

Emerson, K., Nabatchi, T., & Balogh, S. (2011). An integrative framework for collaborative governance. *Journal of Public Administration Research and Theory*, *22*, 1–29. http://dx.doi.org/10.1093/jopart/mur011.

Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). The KDD process for extracting useful knowledge from volumes of data. *Communications of the ACM*, *36*, 27–34.

Fixsen, D. L., Blase, K. A., Naoom, S. F., & Wallace, F. (2009). Core implementation components. *Research on Social Work Practice*, *19*(5), 531–540. http://dx.doi.org/10.1177/1049731509335549.

Flaherty, C., Collins-Camargo, C., & Lee, E. (2008). Privatization of child welfare services: Lessons learned from experienced states regarding site readiness assessment and planning. *Children and Youth Services Review*, *30*, 809–820. http://dx.doi.org/10.1016/j.childyouth.2007.12.009.

Hand, D., Mannila, H., & Smyth, P. (2001). *Principles of data mining.* Cambridge, MA: MIT Press.

Head, B. W. (2008). Three lenses of evidence-based policy. *Australian Journal of Public Administration*, *67*, 1–11. http://dx.doi.org/10.1111/j.1467-8500.2007.00564.x.

Heinrich, C. J. (2002). Outcomes-based performance management in the public sector: Implications for government accountability and effectiveness. *Public Administration Review*, *62*, 712–725. http://dx.doi.org/10.1111/1540-6210.00253.

Heinrich, C. J. (2007). Evidence-based policy and performance management: Challenges and prospects in two parallel movements. *American Review of Public Administration*, *37*, 255–277. http://dx.doi.org/10.1177/0275074007301957.

Hill, C. J., & Lynn, L. E., Jr. (2004). Governance and public management: An introduction. *Journal of Policy Analysis and Management*, *23*, 3–11.

Jennings, E. T., Jr., & Hall, J. L. (2004). Evidence-based practice and the use of information in state agency decision making. *Journal of Public Administration Research*, *22*, 245–266. http://dx.doi.org/10.1093/jopart/mur040.

Kovner, A. R., Elton, J. J., & Billings, J. (2000). Evidence-based management. *Frontiers of Health Services Management*, *16*, 3–46.

Kum, H. -C., Duncan, D., & Wang, W. (2004b). *Understanding social welfare service patterns using sequential analysis: Proceedings of the Fourth National Conference on Digital Government Research, Article 131.* (Retrieved from http://dl.acm.org/citation.cfm?id=112422).

Kum, H. -C., Krishnamurthy, A., Machanavajjhala, A., Reiter, M., & Ahalt, S. (2014). Privacy preserving interactive record linkage (PPIRL). *Journal of the American Medical Informatics Association*, *21*, 212–220.

Kum, H. -C., & Ahalt, S. (2013, March). Privacy by design: Understanding data access models for secondary data. *Symposium conducted at the American Medical Informatics Association (AMIA) joint summits on translation science: clinical research informatics, San Francisco, CA.*

Kum, H. -C., Duncan, D. F., & Stewart, C. J. (2009). Supporting self-evaluation in local government via KDD. *Government Information Quarterly*, *26*, 295–304. http://dx.doi.org/10.1016/j.giq.2008.12.009.

Kum, H. -C., Duncan, D., Flair, K., & Wang, W. (2003). *Social welfare program administration and evaluation and policy analysis using knowledge discovery and data mining (KDD) on administrative data: Proceedings of the Third National Conference on Digital Government Research, 1–6.* (Retrieved from http://dl.acm.org/citation.cfm?id=1123290).

Kum, H. -C., Duncan, D., Flair, K., & Wang, W. (2004a). *Successfully adopting IT for social welfare program management: Proceedings of the Fourth National Conference on Digital Government Research, Article 121.* (Retrieved from http://dl.acm.org/citation.cfm?id=1124312).

Lane, J. (2005). *Optimizing the use of micro-data: An overview of the issues presented at I Quality and Access to Federal Data: Memorial session in honor of Pat J. Doyle: Proceedings of the 2005 American Statistical Association Conference.* (Retrieved from http://www.amstat.org/committees/ethics/linksdir/Jsm2005Lane.pdf).

Lawrence, C. N., Rosanbalm, K., & Dodge, K. (2011). Multiple response system: Evaluation of policy change in North Carolina's child welfare system. *Children and Youth Services Review, 33*, 2355–2365.

Lowndes, V., & Skelchner, C. (1998). The dynamics of multi-organizational partnerships: An analysis of changing modes of governance. *Public Administration*, *76*, 313–333. http://dx.doi.org/10.1111/1467-9299.00103.

Macomber, J., Keuhn, D., McDaniel, M., Vicker, T., Pergamit, M., Cuccaro-Alamin, S., ... Barth, R. P. (2008). *Coming of age: Employment outcomes for youth who age out of foster care through their middle twenties.* Washington, DC: Department of Health and Human Services.

Maynard, R. A. (2006). Presidential address: evidence-based decision making: What will it take for the decision makers to care? *Journal of Policy Analysis and Management*, *25*, 249–265. http://dx.doi.org/10.1002/pam.20169.

McBeath, B., Briggs, H. E., & Aisenberg, E. (2009). The role of child welfare managers in promoting agency performance through experimentation. *Children and Youth Services Review*, *31*, 112–118. http://dx.doi.org/10.1016/j.childyouth.2008.06.004.

Meezan, W., & McBeath, B. (2011). Moving toward performance-based, managed care contracting in child welfare: Perspectives on staffing, financial management, and information technology. *Administration in Social Work*, *35*, 180–206.

Meier, K. J., & O'Toole, L. J., Jr. (2009). The proverbs of new public management. *American Review of Public Administration*, *39*, 4–22.

Millett, L., Lanier, P., & Drake, B. (2011). Are economic trends associated with child maltreatment? Preliminary results from the recent recession using state level data. *Children and Youth Services Review*, *33*, 1280–1287.

Naccarato, T. (2010). Child welfare informatics: A proposed subspeciality for social work. *Children and Youth Services Review*, *32*, 1729–1734.

Nguyen, L. H. (2007). Child welfare informatics: A new definition for an established practice. *Social Work*, *52*, 361–363. http://dx.doi.org/10.1093/sw/52.4.361.

North Carolina Department of Social Services (NC-DSS) (2011). *REAP: NC's new approach to delivering technical assistance.* Raleigh, NC: Author (Retrieved from http://www.ncdhhs.gov/dss/mrs/docs/newsletter_may_2011.pdf).

Penuel, W. R., & Means, B. (2011). Using large-scale databases in evaluation: Advances, opportunities, and challenges. *American Journal of Evaluation*, *32*, 118–133.

Poertner, J., Moore, T., & McDonald, T. P. (2008). Managing for outcomes: The selection of sets of outcome measures. *Administration in Social Work*, *32*, 5–22. http://dx.doi.org/10.1080/03643100802293808.

Salamon, L. M. (2001). The new governance and the tools of public action: An introduction. *Fordham Urban Law Journal*, *28*, 1611–1674.

Schorr, E., & Auspos, P. (2003). Usable information and what works: Building a broader and deeper knowledge base. *Journal of Policy Analysis and Management*, *22*, 669–676. http://dx.doi.org/10.1002/pam.10162.

Sim, I., Gorman, P., Greenes, R. A., Haynes, R. B., Kaplan, B., Lehmann, H., & Tang, P. C. (2001). Clinical decision support systems for the practice of evidence-based medicine. *Journal of the American Medical Informatics Association*, *8*, 527–534.

Stewart, C. J., & Duncan, D. F. (2013). *Changes and trends in the child welfare caseload in North Carolina: April 2013.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina at Chapel Hill.

Stewart, C. J., Kum, H. -C., Barth, R. P., & Duncan, D. F. (2014). Former foster youth: Employment outcomes up to age 30. *Children and Youth Services Review*, *36*, 220–229.

Usher, C. L., Locklin, E., Wildfire, J. B., & Harris, C. C. (2001). Child welfare performance ratings. *Administration in Social Work*, *25*, 35–51. http://dx.doi.org/10.1300/J147v25n01_03.

Usher, C. L., Wildfire, J. B., Brown, E. L., Duncan, D. F., Meier, A., Salmon, M. A., ... Gogan, H. C. (2002). *Evaluation of the Title IV-E waiver demonstration in North Carolina.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina.

Wildfire, J. B., Duncan, D., Brown, E. L., Calix, A., Gogan, H., Green, R., ... Stewart, C. (2007). *North Carolina Title IV-E waiver demonstration termination report.* Chapel Hill, NC: Jordan Institute for Families, University of North Carolina.