Machine Learning in Workforce Development Research: Lessons and Opportunities

Siobhan Mills De La Rosa, Nathan Greenstein, Deena Schwartz, and Charlotte Lloyd ● November 2021

Highlights

This brief summarizes Abt Associates' experiences and lessons from the Descriptive and Analytical Career Pathways Project's machine learning study. These experiences suggest that machine learning:

- Can be a powerful tool in the right context.
- Involves some risk and users should be cognizant of the limitations and expected results of this approach.
- May struggle to replicate the detail or nuance of human research in the context of implementation research.
- May require human researchers to dedicate substantial time and resources to define key concepts.
- May require substantial input from human researchers.
- May require a team with interdisciplinary skill sets to be completed successfully.
- Operates in an evolving legal, computing, and cost environment.

Government entities and social science researchers are increasingly interested in using machine learning to shed light on important policy questions. ***Machine learning*** refers to a range of computer models and algorithms able to uncover patterns, create categories, and make predictions from large data sets without being given step-by-step directions from a human (Alpaydin, 2009). Machine learning methods allow researchers to analyze a larger volume and wider breadth of data than could be reasonably be done by humans. Machine learning can be used to identify and strengthen insights on a range of issues such as public perceptions of key topics, products, or services; themes in program implementation or service use; and predictions of users' future behavior based on current and historical data.

As part of the Career Pathways Descriptive and Analytical Project (CP D&A) sponsored by the U.S. Department of Labor (DOL) and conducted by Abt Associates, this brief presents findings from an exploratory study examining how machine learning can be used to synthesize a large body of data about the implementation of career pathways programs. Career pathways programs aim to help individuals enter and exit occupational training at different levels—depending on their initial skills and work experience—and advance over time to higher skills, industry-recognized credentials, and better jobs with higher pay (see *The Career Pathways Approach* box). These programs generally involve a wide range of funding sources, service delivery systems, and target populations (Fein, 2012). In the last decade, interest in and evaluation of career pathways approaches have grown dramatically, including at DOL (e.g., Sarna & Adam, 2020). The study: (1) explores how machine learning can be used synthesize and draw lessons from available text-based data to provide comprehensive information on implementation career pathways programs and (2) provides lessons on how machine learning could be used in future workforce development research.

Chief Evaluation Office
U.S. DEPARTMENT OF LABOR

About the Study

The Workforce Innovation Opportunity Act (WIOA) emphasizes the use of career pathways programs and requires the Department of Labor (DOL) to conduct a study to develop, implement, and build upon career advancement models and practices. In order to respond to the need for information and evidence in the field due to this growing emphasis, DOL's Chief Evaluation Office, in collaboration with the Employment and Training Administration, contracted with Abt Associates to conduct the Descriptive & Analytical Career Pathways Project. The project's purpose is to advance the evidence base in the career pathways field by addressing key research gaps, drawing primarily on existing data, to inform career pathways systems and program development to help meet the needs of both participants and employers.

The Career Pathways Approach

The approach involves a combination of rigorous and high-quality education, training, and other services (WIOA, 2014) and has four main tenets (e.g., Fein, 2012; Werner et al., 2013): The career pathways approach:

- offers articulated steps in an industry sector, offering multiple places to enter and exit training;
- results in recognized credentials that intend to lead to better jobs with higher pay;
- uses support services and provides flexibility needed for non-traditional students; and
- relies on employer connections and partnerships.

The machine learning study of career pathways programs' implementation was designed to address challenges in synthesizing the existing and large body of information. No centralized data source is available that provides comprehensive and consistent information on career pathways program implementation. Moreover, while over 80 studies have examined implementation of this approach (Sarna & Adam, 2020), many more have not been evaluated formally. A growing body of web-based text data on career pathways programs is available from program websites, including descriptions of their programs, training offerings, and student services. Machine learning methods provide an opportunity to analyze these text data, which have not traditionally been analyzed in large volumes because accessing and analyzing them at scale is difficult.

This exploratory study was designed to explore how to use machine learning methods, including web scraping, supervised learning, and natural language processing, to collect and analyze a large volume of data on career pathways program implementation. The study design included two phases: (1) an initial data collection phase to identify an approach and techniques for collecting data, and (2) an analysis phase to synthesize the data. While only the first phase was completed as part of this project, the experience yielded important lessons for future machine learning studies.

This brief summarizes lessons learned from using machine learning to study the implementation of career pathways programs. First, this brief first describes the research questions that guided the study and summarizes the machine learning methods designed for the data collection and analysis activities, including study limitations and challenges encountered. It then provides lessons learned on using machine learning methods for social science research. Finally, the brief discusses strategies for using these methods in future workforce development projects and other areas, particularly federally funded efforts.

Key Terms in Machine Learning

Web scraping, also known as ***scraping,*** refers to the process of collecting data from websites using machine learning tools.

Supervised learning algorithms are machine learning algorithms that depend on a "training" set of data that has been manually labeled by researchers to indicate the true or "correct" action, decision, or other output for each observation in the data set. A supervised learning algorithm iteratively processes these data with the goal of "learning" to replicate the true outputs it was given. The algorithm can then construct outputs for data it has never seen before.

Natural language processing (NLP) is a family of techniques designed to extract meaning from unstructured human language, including text documents, handwriting, and speech, using computational and linguistic models.

1. Machine Learning Study: Research Questions and Methods

A. Research Questions and Study Limitations

The primary goal of the CP D&A project's machine learning study was to identify career pathways implementation lessons from a large volume of implementation data and to explore the potential of machine learning methods for workforce development research. The study aimed to answer two sets of research questions, one focused on career pathways implementation and one focused on the potential of machine learning as research methods.

Specifically, the first set of questions examine several dimensions of career pathways implementation:

- How are career pathways programs and systems being described and implemented? What components or elements do they include?
- To what extent are these career pathways programs and systems implementing each of the elements described in the Workforce Innovation and Opportunity Act (WIOA) (for program-level efforts) or Six Elements[1] definitions (for systems-level efforts)?
- What themes arise from career pathways program and systems descriptions that are *not* reflected in the WIOA and Six Elements definitions?
- To what extent are these career pathways programs implementing *multiple* elements of the career pathways model in combination with each other? And to what extent are career pathways efforts comprehensive in their implementation of career pathways models?
- What is the prevalence of systems-based career pathways efforts? In what ways do systems-based efforts differ from individual-program-based efforts?

For these research questions, the study was designed to capture which career pathways program components are most frequently described on webpages. Because machine learning algorithms can only identify patterns that are explicitly present in the data and because they generate findings from a wide range of data, the results are not as likely to be impacted by researcher biases about what should be appearing in the data. In addition, it can identify themes of implementation features that had not been previously observed in the literature. For example, one level of insight would have been comparing these most frequently documented features to the key program components cited in both the WIOA and Six Elements definitions. An additional level of insight would have been to identify themes of implementation features that had not been previously observed in the literature.[2]

1 DOL's Six Elements refer to the six elements that are necessary to developing a comprehensive career pathways program. These are: 1. Build cross-agency partnerships and clarify roles. 2. Identify industry sectors and engage employers. 3. Design education and training programs. 4. Identify funding needs and sources. 5. Align policies and programs. 6. Measure system change and performance.

2 Because machine learning allows for collecting and analyzing data from a larger number of career pathways programs, the study could potentially capture features from these programs that are not yet examined in the existing literature. Second, because machine learning generates findings that are empirically based, the results are not as likely to be impacted by researcher biases about what should be appearing in the data. Word and term frequency findings may suggest some concepts that have not been highlighted in career pathways frameworks are appearing in the data that had not previously been identified.

The second set of questions aimed to learn about machine learning's appropriateness as a tool for career pathways research.

1. What can machine learning tell us about career pathways? What can it not tell us?
2. What are the strengths and weaknesses of using machine learning as an analysis tool in social science and government-contracted research on career pathways or in similar contexts?
3. What data sources did we draw on using machine learning?
4. To what extent can we reach new conclusions by drawing on the much larger body of data that machine learning allows us to harness?

While machine learning is a powerful tool to analyze large amounts of data, the approach has some limitations. These are discussed in more detail throughout the brief, but it is important to recognize that machine learning can only analyze patterns that exist in the data. Practically speaking, this means that the implementation features identified could only include those features that programs described and may not include particular components of interest to program administrators and policymakers. In addition, while the word and term frequencies that natural language processing algorithms generate can provide useful insights on what is contained in a large dataset, they may not provide the kind of detailed or nuanced information that may be important to policymakers and program administrators. For example, while machine learning can potentially identify the most commonly implemented strategies, it may provide more limited information on particular topics of interest such as service sequencing and approaches to address particular challenges, such as low completion rates.

B. Machine Learning Data Collection and Analysis Activities and Study Challenges

Developed in consultation with DOL, the study used a two-phased plan consisting first of data collection and then analysis:

- The study first designed a plan to collect text data focused on career pathways implementation from a wide range of web-based sources. These included the existing research literature, program descriptions written by career pathways programs and providers themselves, as well other sources of career pathways-relevant information.
- The study then planned to use a combination of analytical techniques to identify implementation themes across a large and diverse body of data. However, as discussed in more detail below, the second phase of the original design was not completed.

Exhibit 1 provides an overview of the steps in the planned data collection and data analysis processes, and each is discussed in turn below.

Exhibit 1: Data Collection and Corpora Building Process for the CP D&A Machine Learning Study

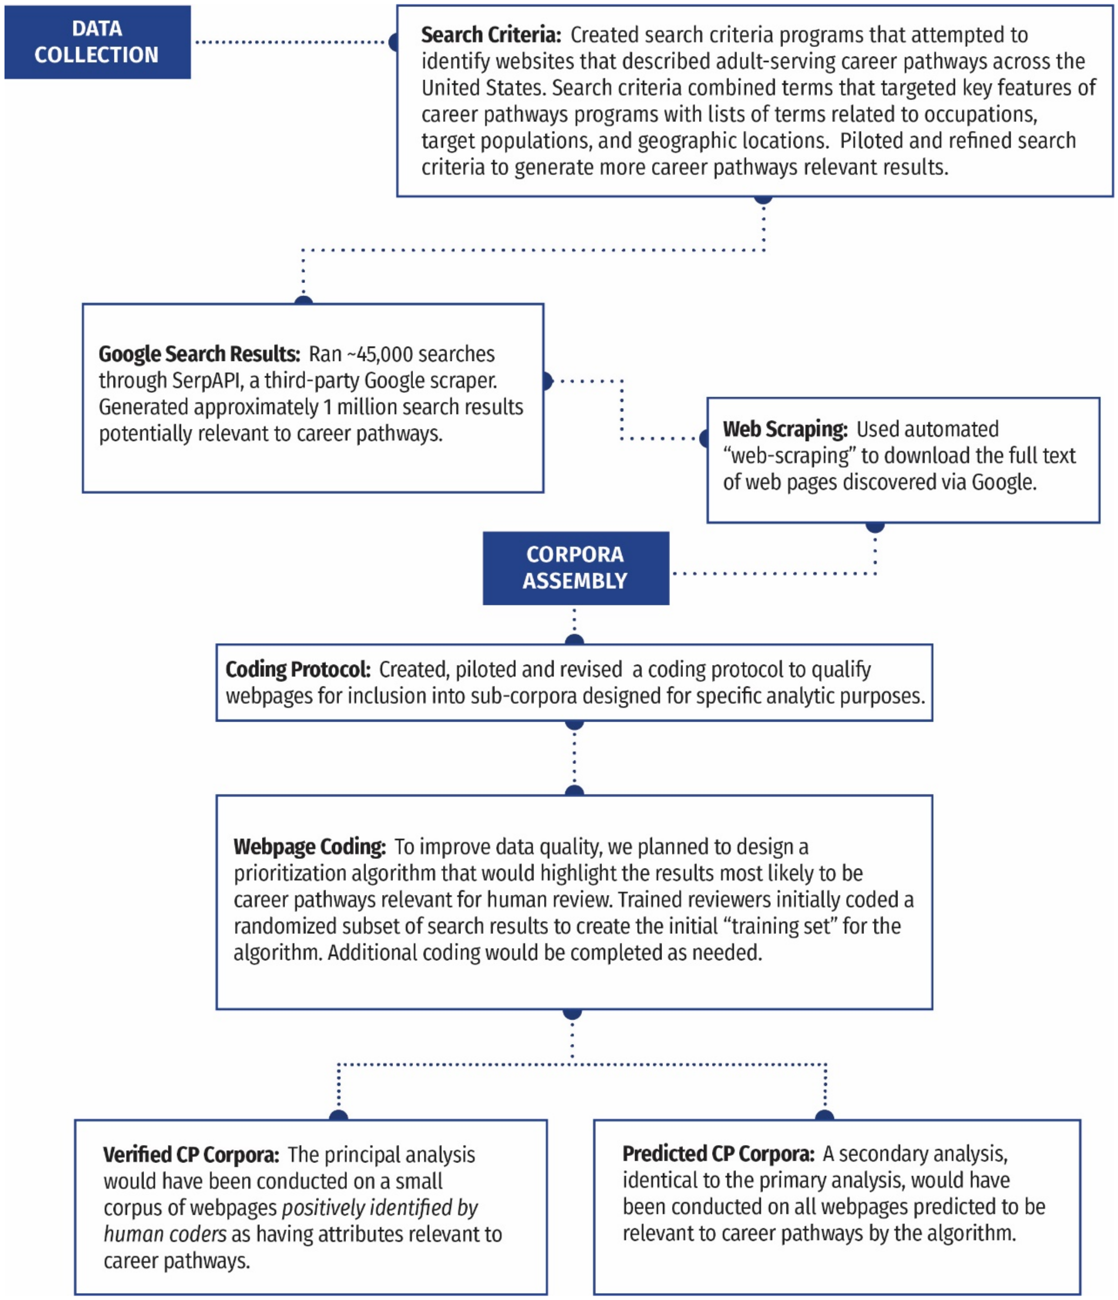

Phase 1: Machine Learning Data Collection Activities

A ***corpus (plural: corpora)*** refers to a collection of texts organized such that they can be analyzed easily. For this study, the corpora essentially functioned as the analytic datasets.

The first phase of the study focused on identifying potential career pathways-relevant data sources, scraping text data from these sources, and assembling those data into various datasets, known as corpora (see box). Based on the research questions, the study team aimed to collect text data that captured how career pathways program operators (community colleges, non-profits, workforce development boards, etc.) described themselves and their programs, in part to understand how actual implementation may differ from what is discussed in research and legislation, or by funders (see Sarna & Strawn, 2018). Exhibit 2 provides an overview of the types of data the study team collected on career pathway programs to address the research questions.

Exhibit 2: Data for the CP D&A Machine Learning Study

Data Collected	Source	Provides Information On…
Descriptions of career pathways program implementation written by independent authors.	Existing research literature, reports from funding agencies, and news articles.	Program implementation from a third-party observer.
Descriptions of career pathways program implementation written by the program staff.	Program descriptions from career pathways program websites and annual reports.	Program implementation from program operators. Allows researchers to determine what components of career pathways are most frequently implemented, according to the people who implement them.
General information about career pathways programs, models, and frameworks.	Local, state, and federal government websites; websites of philanthropic institutions and think tanks; community college websites.	Career pathways programs generally, including their motivations and structure.
Descriptions of career pathways efforts at the systems-level.	Websites of state and federal governments as well as philanthropic entities and State WIOA plans.	Efforts to create and promote career pathways initiatives across a system of education, workforce development, or non-profit providers.

To collect the data, the study developed a detailed, transparent, and reproducible project definition of “career pathways” with which machine learning algorithms could work to guide the overall web search. Although clear definitions of career pathways programs have been established, machine learning algorithms are most effective with definitions that are narrowly defined so that the algorithm easily “identifies” career pathways programs in the text.[3] The study used the existing literature and feedback from subject matter experts to “operationalize” an algorithm-friendly definition of career pathways that broke the concept into its six key components:

- Offers education or training for one or more specific occupations in a specific sector or industry;
- Is not exclusively targeted to high school students;
- Does not require a Bachelor’s or Associate degree for entry;
- Provides education or training that results in a credential (i.e., certificate, technical diploma, degree, certification, license);
- Indicates how the coursework or credential contributes to later credentials (i.e., pathway, ladder, lattice), either offering subsequent credentials as part of the program or clearly indicating the next step or steps in credentialing. The credential associated with the first step must be below a Bachelor’s degree but more than a day-long training (definition excluded very short-term credentials such as CPR, ServSafe, OSHA-10); and
- Offers individualized academic, career, and/or logistical (e.g., transportation, childcare, financial planning) supports. In a post-secondary institution setting, career pathways students must receive support services beyond what is available to students not in a career pathways program (i.e., services beyond a career office, tutoring center, or academic advising on campus generally).

This operationalized definition sometimes includes more specific criteria than those used by career pathways program operators when describing their programs. However, defining career pathways in this way allowed the study to identify programs that, as described, implemented all key career pathways components.

Based on this operationalized definition, the study then created a robust set of Google search terms that were used to conduct the Internet search. Using search strings that combined key career pathways terms with location, target population, and occupation search terms, the study designed strings to be inclusive to capture the largest number of relevant results.[4] A third-party application programming interface (API) was used to conduct 44,281

[3] Unlike human researchers, machine learning algorithms cannot independently apply contextual knowledge or subject matter expertise when determining how data sources should be categorized. As such, the “definition” for the purposes of this study had to break career pathways down into its most critical components so that the algorithms could identify them in the text data.

[4] Programs that exclusively served youth under the age of 18 were not included.

independent Google searches. The full set of searches yielded approximately one million unique search results, which were used for the web scraping.

As planned, the raw text data collected during the scrape would have been turned into a series of corpora that could be used for analysis. To produce these corpora, the study piloted and refined a coding protocol to tag webpages for inclusion in particular corpora for analysis. As described in Exhibit 3, the study defined six mutually exclusive categories into which webpages could be filtered, plus a seventh "systems" category that could be applied to results coded as "CPI-1," "CPI-2," and "CP" if the results represented systems-level efforts. Webpages in the Career ("C"), Noise ("N"), and Broken ("B") categories would not be included in the analysis.

Exhibit 3: Coding Protocol Values and Intended Usage for the CP D&A Machine Learning Study

Label	Pages in this Category…	Would Answer Questions About…
Career Pathways Initiative – 1 (CPI-1)	Included descriptions of career pathways programs that were written by the program's staff. For example, a webpage that described the Heating, Ventilation, and Air Conditioning (HVAC) technician career pathways program at a local community college, published by the community college.	• How do program operators describe their program? • What are key features of career pathways programs as implemented?
Career Pathways Initiative – 2 (CPI-2)	Included descriptions of career pathways programs that were written by someone other than the program operators. For example, an independent implementation study that described the program components of career pathways programs at a community college.	• How do independent observers describe career pathways programs? Do these descriptions differ from the descriptions in CPI-1? • What are key features of career pathways programs according to independent observers?
Career Pathways (CP)	Did not describe a particular career pathways program, but did include information about career pathways programs more generally. For example, a webpage describing DOL's Trade Adjustment Assistance Community College and Career Training (TAACCCT) initiative.	• What are high-level features of career pathways programs? • How do these descriptions vary from those in CPI-1 and CPI-2?
Career (C)	Included information that was relevant to a career in a specific occupation, but did not include any career pathways-relevant information.	Results tagged as C were not included in analysis.
Noise (N)	Did not include career pathways-specific information but were picked up by the search results for other reasons.	Results tagged as N were not included in analysis.
Broken (B)	Had broken links at the time of review. A larger percentage of links were broken at the time of review than expected, suggesting that the scrape and review of search results should happen as soon after the search as possible.	Results tagged as B were not included in analysis.

In preparation for the supervised learning step, human coders reviewed a randomly-selected subset of the scraped search results and flagged each into one of the categories in Exhibit 3.

Phase 2: Machine Learning Analysis Activities

For its second phase, the study planned to analyze those data using a variety of *natural language processing* algorithms (Kurdi, 2016) to distill common themes from the text data. These included:

- **Bag of words** analyses to produce word counts that reveal which words appear most frequently within and across documents, yielding insights into which themes are most important within the dataset (Goldberg, 2017, p. 69).

- **N-grams** to quantify the frequency of two, three, or n-term phrases within a corpus by looking for sequences of contiguous words. Given the variety of terms used to describe similar career pathways features, n-grams would have helped us empirically identify which terms the field is most commonly using to describe itself (Goldberg, 2017).

- **Topic modeling** to identify more abstract concepts represented in a corpus (Blei, 2012). The advantage of topic modeling is that it helps researchers determine potential abstract topics (not just words or phrases) that are common within a corpus of documents (Blei, Ng, & Jordan, 2003).

As designed, the next step in the analysis phase would have been to develop a machine learning prioritization algorithm to identify the text data most likely to be relevant to career pathways. This would have consisted of manually coding search results until enough were coded to train the algorithm. The algorithm would have been trained to predict the codes most relevant to all remaining scraped search results not yet coded manually. Ideally, these predictions would funnel promising data to human reviewers and identify the data not reviewed by humans that was most likely to be relevant to the analysis (see Exhibit 1 above). Once this step was completed, the study would have created the analytic corpora and begun the natural language processing analysis.

Key Study Challenges

As discussed, the second phase of machine learning study was not completed as part of the CP D&A project. This occurred as the result of several factors that made the scope of the work difficult to complete within the parameters of the CP D&A project.

- *Data collection efforts yielded a much larger volume of data than anticipated.* During the design phase, the study anticipated identifying several hundred websites that would contain career pathways implementation information. In practice, our search terms generated approximately one million Google search results. This increased the scope of the project beyond what was designed.
- *Many webpages were only tangentially relevant to career pathways.* A review of a randomly selected subset of the webpages suggested that many were not about career pathways programs. Conducting the planned analysis on a dataset that did not remove the "noisy" data would have produced findings that did not represent career pathways implementation. However, removing the noisy data would have required a much more sophisticated and resource-intensive approach to data collection and corpora assembly that was beyond the scope of this project.
- *Storing and analyzing larger volumes of data required more significant computing resources than were available in a typical computing environment.* In order to operationalize the planned prioritization algorithm, the study would have faced additional costs in terms of staff time needed to build the cloud computing environment, integrate it into existing systems, and provide ongoing oversight and server maintenance.
- *Extensive consultations with legal counsel were required after unanticipated legal questions arose.* The laws and legal precedent around machine learning methods is rapidly evolving. Consultations with legal counsel were required to determine whether 1) websites' terms of services prohibited web-scraping, 2) the study could use a third-party vendor to scrape Google results, and 3) the study could use particular analytical tools that had prohibitions around use for profit-making activities.
- *Uncertainty in the methodology's ability to identify programmatically relevant findings with available resources.* As discussed, a limitation of the study was that analyzing text using machine learning can produce results that may be too general to be relevant to some policymakers and practitioners. These uncertainties were difficult to address given the goals and resources for the CP D&A project.

In spite of, and also because of these challenges, the experience using machine learning on the CP D&A project yield several important lessons for future initiatives. These are discussed in the next section.

2. Lessons on Using Machine Learning Methods in Workforce Development Research

Based on the CP D&A experience, as well as Abt's subject matter and machine learning expertise, this section highlights seven important lessons about using machine learning methods in research focused on program implementation. The discussion focuses on machine learning as it relates to DOL's needs and interests around workforce development policies and programs. Although the lessons are relevant to a broad range of potential machine learning methods and efforts, they are most applicable to future social science projects with goals and approaches similar to this study.

MACHINE LEARNING CAN BE A POWERFUL TOOL IN THE RIGHT CONTEXT.

The term "machine learning" covers a wide array of analytic approaches and tools which must be appropriately tailored to the research questions of interest and the data that will be analyzed. Based on the experience of the CP D&A study, machine learning methods ***may be useful*** when researchers or other stakeholders:

- *Have a clear set of research questions that lend themselves to machine learning or cannot feasibly be answered using traditional methods.*
- *Are working with a known dataset that they understand or have sufficient time and resources to explore.*
- *Are interested in predicting future behavior and have the data to build a training data set.*
- *Want to examine patterns in data that are difficult for human researchers on their own.*
- *Wish to expedite a data processing task.*
- *Want to conduct exploratory analyses for new research, either with known data or with a manageable amount of data that is not well understood by researchers and stakeholders.*
- *Are interested in bringing an unbiased and primarily data-driven strategy into an analysis.*[5]

Machine learning has advantages and disadvantages (Nadkrani et al., 2011), and the appropriateness of the technique depends on the nature of the research question to be addressed. One advantage of machine learning is that it can be used to process and analyze large volumes of data – much more than could be analyzed using traditional means. It can also be used to identity underlying patterns in data that may be difficult for human researchers to detect. Under certain conditions, machine learning can be used to predict the behavior of individuals, markets, and other phenomena. Overall, machine learning methods allow researchers to answer questions that would have been challenging or infeasible to answer using traditional methods.

Machine learning methods present some drawbacks, however. Researchers using machine learning methods can lose nuance and the ability to conduct in-depth analyses in target areas. For example, natural language processing algorithms are not designed to provide the contextualized and operational insights that program practitioners may find helpful, such as the intensity of program staff contacts with students needed to improve attendance or how support services could be structured to help students overcome barriers to participation and employment. Analyzing a large volume of data with machine learning also means that researchers may not be able to review their data exhaustively. For example, analyzing text data from thousands of websites may limit researchers' ability to view the content of each individual website and determine its relevance to the research effort. Additionally, although researchers may know the likely form of results they will produce from traditional analyses (i.e., table shells, a final report, etc.), machine learning methods may produce results that are harder to understand without additional explanation (e.g., topic models). Practitioners may find that results from machine

---

5 Machine learning, when combined with theory-driven analysis, can offer a way to reduce the influence of researcher biases, preconceptions, or blind spots (Holzinger & Jurisica, 2014; Berdanier et al., 2020).

learning may not be as immediately and easily applied in their work (e.g., findings that suggest potential bias or areas for additional inquiry) (Berdanier et al., 2020).[6]

Based on the experience of this project, the level of detail with which any machine learning effort can answer any specific question of interest will depend on how well-defined the questions and topics of interest are at the outset, the data that researchers are using, and the machine learning methods chosen. Based on the CP D&A experience, pilot testing and exploratory research can lend clarity to this picture before the research process begins, but an iterative process is likely needed. Collaboration between subject matter experts and "data scientists" (researchers with extensive experience using machine learning methods) appear to be important in navigating these and other tradeoffs of machine learning research.

Based on this project's experience, machine learning methods ***may not be useful*** in addressing a research question or topic when the research or project:

- *Requires detailed or nuanced answers about loosely defined concepts*. If research topics are not uniformly defined or guided by a strict set of rules, many machine learning algorithms may not imitate the judgement and output of human subject matter experts closely enough to address DOL's research needs. Training and tuning an algorithm to expert-level nuance might require resources beyond what government contracts can reasonably provide.
- *Relies on data that is low-quality or unlikely to contain information on the patterns stakeholders are trying to identify*. This may be because the available data do not capture an influential third factor. To take an extreme example, even a highly sophisticated machine learning model will be unable to use unemployment insurance claims data to predict filers' favorite flavors of ice cream. A model trained on grocery store purchase data, however, could likely perform this task well. As in most research, success depends on the data available and how closely they relate to the outcomes of interest.
- *Requires an extensive collection of data from sources that are not well understood or have not been explored by researchers*. Researchers using large volumes of un-vetted data will require significant resources to understand the dataset and design approaches that can appropriately analyze it.
- *Has time, resource, and/or other constraints that do not allow flexibility for the exploration and iteration that machine learning entails*. Machine learning efforts require a higher degree of flexibility and iteration than traditional research projects. Creating contracts with additional flexibility built into machine learning tasks can ensure the success of these projects.
- *Could use traditional methods to achieve similar or better results for similar or fewer resources*. For example, when a traditional quantitative or qualitative research design would satisfy project needs, using machine learning methods may not be an appropriate use of resources.

In sum, based on the CP D&A experience, the value of machine learning methods is that, under the right circumstances, they can allow researchers to answer questions using data that could not be easily analyzed before or were too costly to analyze (text data, satellite imagery, location histories, web-browsing habits, etc.). They can also uncover new patterns in existing data and identify patterns in data that were not easy for human researchers and traditional statistical tools to detect such data bias.

[6] Machine learning methods will highlight the relationships that are mathematically strongest in the data. As a result, these algorithms will sometimes produce findings that are not obviously policy actionable. For example, topic modeling can identify words or phrases as being especially important, but that are difficult to map onto specific implementation characteristics, such as recruitment language intended to advertise programs to students. In such a case, findings may be challenging to interpret, but they can yield unexpected insight and inspire new ways to think about program design and implementation. In general, it may be difficult to anticipate what relationships machine learning algorithms will identify as mathematically strongest. Researchers may encounter relationships that they were not anticipating, which requires further research into these relationships to provide context and develop further understanding of its application.

ANY GIVEN APPLICATION OF MACHINE LEARNING METHODS INVOLVES SOME RISK, REQUIRING AN UNDERSTANDING OF THEIR LIMITATIONS AND EXPECTED RESULTS.

Machine learning methods have been incorporated into the day-to-day work of fields such as information technology (IT) and medicine, but they have not been widely applied in workforce development or other social science research. Based on the experience of the CP D&A project, work is ongoing to determine how these methods can address important questions in workforce development research. However, the risks and limitations of using machine learning methods should be recognized and may include:

- *More "noisy" data than originally anticipated.* The study found that a primary challenge for this project was identifying the small number of websites that had career pathways implementation information out of the universe of websites. As discussed, this challenge was more substantial than the study anticipated during the design phase. To identify websites with relevant career pathways information, the study created sophisticated search queries that would identify as many relevant websites as possible while filtering out websites that were not relevant (i.e., "noise"). IT firms have done this successfully for workforce development projects (Agrawal et al., 2017), but this filtering effort takes substantial resources, as well as a close working relationship between programmers and human reviewers with substantial subject matter expertise. Though projects using a "scrape the Internet" approach can be successful, they will require substantial resources and IT capabilities.
- *Unstructured data that require substantial human review.* In the CP D&A experience, working with large volumes of unstructured data from the Internet may require a greater level of human review of data quality than is feasible. Text data are often unstructured and unpredictable, especially when drawn from unknown or varying sources (e.g., scraped from across the Internet). The large volume of data often can make it impractical or cost-prohibitive for researchers to individually review every data source. Particularly for web-scraped data that is likely to have uncertain content, researchers may elect to identify creative approaches that balance reviewing data to improve quality and available resources. Researchers can opt to review every data source, which can improve data quality, but, as a result, they may need to limit the scope of the analysis to the amount of data that can be manually reviewed given finite project resources. In such cases, machine learning can still be used to help human reviewers work more efficiently than they otherwise could. Researchers with resource constraints may need to navigate tradeoffs between breadth, resources, and data quality regardless of whether they employ traditional or machine learning methods. Still, the increased volume of data involved in machine learning practices can make these tradeoffs more salient than with traditional methods.
- *Substantially more iteration than is typical in traditional analyses.* Iterating through multiple rounds of data collection, cleaning, and analysis, which is typical in machine learning projects, can consume significant time and resources. As in any research project, researchers using machine learning can expect to iterate through their analysis several times to refine and revise their analytical approach. However, because supervised machine learning algorithms "learn" by picking up patterns in training data sets, machine learning projects will often iterate through an analysis many more times than a standard data analysis project as the algorithm learns and improves its performance.

IN THE CONTEXT OF IMPLEMENTATION RESEARCH, MACHINE LEARNING METHODS MAY HAVE DIFFICULTY REPLICATING THE DETAIL AND NUANCE CAPABLE OF HUMAN RESEARCHERS.

While machine learning algorithms' ability to process large quantities of data – known as *big data* – is important in some contexts, machine learning algorithms will struggle to match the nuance of analyses by human researchers. For example, a human research team can read through a subset of documents and determine the key components of a program, the sequencing of those components, and how any contextual factors may contribute to program implementation. A machine learning algorithm cannot

Big data describes collections of information that until recently would have been unmanageably vast or complex to analyze. Examples include electronic health records, satellite imagery, and text from social media posts.

Training data sets refer to datasets that are used to "teach" supervised learning algorithms to identify the pattern of interest to researchers. Researchers typically rely on pre-existing data but could create a training data set from scratch. Larger datasets will give an algorithm more opportunities to "learn" the target pattern.

do this without substantial training from a team of subject matter experts and a large training data set. This is especially true when the algorithm is working with a relatively small volume of data or when the concepts of interest are not clearly defined and require a high level of expertise to categorize accurately.

The natural language processing analysis planned for the CP D&A project may have shed light on important features of career pathways programs, and while useful, it is not clear whether the findings from that analysis would have been detailed enough for use by some practitioners and policymakers. When applied to implementation research, machine learning is likely to afford additional breadth synthesizing larger volumes of information and including previously unknown information rather than depth.

SOME PROJECTS SHOULD EXPECT TO DEDICATE SUBSTANTIAL TIME AND RESOURCES TO DEFINING KEY TERMS, CREATING AND APPLYING A CODING FRAMEWORK, AND INTERPRETING RESULTS.

Lack of standardization around key terms in workforce development can make implementing certain machine learning tasks a challenge. Based on the CP D&A experience, projects applying machine learning methods designed to provide information on concepts without a set of standard definitions and rules may face similar issues.

In this project, career pathways*,* like many key terms in workforce development, is a concept that was defined in overlapping ways by career pathways program implementers, funding agencies, and the research community. In addition, like other workforce development services, career pathways includes "bundling" of different activities (such as occupational training and support services). In order to use machine learning methods to synthesize career pathways implementation information from web sources, the study had to create a narrow definition that broke career pathways into key components that algorithms could easily recognize in the text data. Based on the CP D&A experience, creating, validating, and refining this definition required substantial time and resources, but it was important in moving ahead with this project.

Creating that definition was useful in building the Google search queries to collect relevant data for the CP D&A project. Specifically, the definition helped determine which keywords would be needed to capture both programs that self-identified as career pathways programs and those that did not. The definition represented an "ideal" program based on common features of a variety of career pathways definitions, including the one used by WIOA. Though stricter than what is used by some programs, researchers, and funding agencies, the study used this definition to narrow search results to those most relevant to DOL.

> ***Predictive analytics*** projects aim to make predictions about the unknown by analyzing existing data. A supervised learning algorithm looks for patterns in data that include both predictors and outcomes. Then, given new data, the algorithm uses the patterns it has found to predict new outcomes.

Because the Google searches yielded information on a varied set of programs, the study built the coding protocol to identify which element(s) of the study definition a program did not include. In doing so, the study developed a process to identify which features of the definition were least well reported in the data. As discussed, the purpose of this process was to identify programs that did not meet certain aspects of the study's definition, but the study did not undertake this step.

HUMAN RESEARCHERS ARE NEEDED TO DEFINE THE RESEARCH PROJECT, ITERATIVELY "TEACH" SUPERVISED LEARNING ALGORITHMS TO RECOGNIZE APPROPRIATE PATTERNS, AND INTERPRET RESULTS.

Machine learning algorithms often require substantial human guidance to perform optimally (Alpaydin, 2019). While the algorithms work by recognizing underlying patterns in data, they cannot independently define concepts of interest, conduct analysis, or interpret findings without the involvement of human researchers. Though subject matter experts are often needed on such projects, based on the experience of the CP D&A project, they can guide machine learning projects in four key ways:

- *Subject matter experts can help define key concepts in a transparent and reproducible way.* Every machine learning project's needs will be different, but based on the CP D&A project, machine learning projects involving data collection or supervised learning are likely to need subject matter experts to create narrowly defined definitions that can be operationalized for machine learning algorithms, as this study did for "career pathways" on this project. Although machine learning methods excel at finding solutions that humans may not be able to find, human input is needed to define the problem to be addressed, especially because these algorithms lack the context and subject matter knowledge that human researchers can take for granted.

- *If using supervised learning, subject matter experts can "teach" supervised learning algorithms.* Often, if using supervised learning, subject matter experts must create and apply a coding protocol that can "teach" machine learning algorithms to identify the distinctions of interest. For this project, human coders with subject matter expertise categorized a subset of the results scraped from the web into one of six categories (see Exhibit 2 above). In the planned analysis, this subset would have comprised a stratified sample of all web results, increasing the likelihood that a range of program types and contexts would be represented. The machine learning algorithm would then have "learned" to link the decisions made by our human coders to underlying patterns in the language of the web results; applied this knowledge to results that had not been reviewed by a human, predicting the most appropriate category for each result; and thereby surface those results most likely to be relevant to career pathways.

  Without the initial training dataset created by human coders, the machine learning algorithm would not have been able to predict which websites were most relevant to career pathways. Based on the CP D&A experience, the quality of the training dataset depended on the accuracy of the decisions that human coders made. As such, it is important to use human coders with sufficient subject matter expertise and to provide training on the coding protocol. This project used mid-level analysts with experience in career pathways and workforce development programs to double-code a subset of websites that a senior coder then reconciled.

  While employing knowledgeable human coders can add to the cost of a machine learning project, it is likely to improve the quality of a predictive algorithm trained on a new analytical dataset.[7] The costs to consider include both developing a coding protocol and performing human coding on the training portion of the data.

> ***Supervised learning*** algorithms can open new research possibilities by "learning" from a coded portion of a massive dataset, and then rapidly applying what they have learned to the remainder of the data, with next to no human involvement. "Training" supervised learning algorithms often requires input from subject matter experts and many rounds of iteration. Teams that cannot use a pre-existing training data set and have to build their own will use substantially more resources to do so.

- *Subject matter experts can contextualize the results of machine learning, interpret their meaning, and translate them into operational insights.* Algorithms do not have access to or incorporate important contextual factors in their analysis. Human researchers can situate findings generated by machine learning methods in the appropriate context, interpret their meaning for the field, and translate them into conclusions that are ready to be disseminated and acted upon.

[7] The study collected a larger volume of data than anticipated. In our designed analysis, the study team anticipated collecting a smaller volume of data and planned to review data quality by hand rather than using a prioritization algorithm. When the volume of data became clearer, it was no longer feasible to assess whether all potential data sources were career pathway relevant using human researchers. The prioritization algorithm would have allowed us to train an algorithm to do this, but this was not in the original scope of the project and would included extra costs related to creating a coding protocol, coding a subset of webpages to create the training dataset, building and testing the algorithm, and then deploying it on the full dataset. Given the size of our data, using a prioritization algorithm would have led to increased computing costs to store data in memory for analytic use.

- *Subject matter experts will interact with findings differently depending on the type of machine learning methods employed.* When a project focuses on predictive analytics, a type of machine learning, subject matter experts typically advise on the circumstances under which predictions will be publicly released. Based on the experiences of the CP D&A project and other projects, this can include a schedule for generating predictions, which stakeholders should have access to them, how they should be framed, and what other information should accompany them. From their understanding of the broader systems surrounding the object of prediction, subject matter experts can also help researchers avoid unintended consequences; for example, by envisioning possible future effects, evaluating the potential to do harm, and identifying any blind spots that should be addressed. Finally, subject matter experts can play an important role in helping researchers assess a model's predictions for bias. For example, a predictive analytics project designed to determine who is given access to an Individual Training Account would need to be evaluated for potential bias (i.e., whether the algorithm was giving or denying access more frequently to individuals in certain groups) by subject matter experts before use. Given the many facets of equity and fairness, many predictive analytics efforts require consideration of how potential bias should be evaluated and mitigated, which depends heavily on the context in which the predictions exist.

Unsupervised learning algorithms do not rely on humans to provide examples of distinctions of interest (e.g., a training dataset); instead, these algorithms exploit naturally occurring groupings or relationships in data to reveal insights.

Overall, based on the experience of this project, when a project uses unsupervised learning algorithms or natural language processing, the main role of subject matter experts is guiding data collection and interpreting the findings surfaced by algorithms. Subject matter experts may also need to collaborate with data scientists to evaluate potential data sources for quality and potential bias. Further, machine learning methods yield a range of results, including clusters of apparently related data points, lists of common phrases, or abstract topics appearing in a corpus of text documents. In this project, for example, the clusters and topics generated by machine learning algorithms would not have come with clear labels such as "nursing programs" rather, they would have been expressed as lists of the most relevant text documents or data points, and/or lists of relevant phrases (e.g., "nurse," "Certified Nursing Assistant (CNA)," and "Registered Nurse (RN)").

Subject matter experts also may need to apply their content and context knowledge to the algorithm results, to determine which results are valuable and how they map onto real-world concepts. Having done so, experts can guide more detailed analyses and target comparisons or syntheses of particular data. Finally, experts can link the results of analysis to specific policies and programs that may be useful to practitioners and policymakers.

MACHINE LEARNING TEAMS ARE LIKELY TO BENEFIT FROM INTERDISCIPLINARY SKILLS.

Based on the CP D&A experience, machine learning for policy and program-focused research projects can benefit when staff have complementary skills in four key areas:

- *Project Management.* Given the relative level of uncertainty associated with machine learning projects, project management staff will need to have strong problem-solving and communication skills, especially because they will need to work effectively with a team that has a diverse set of skills and content knowledge.

- *Subject Matter Expertise.* Subject matter experts play an essential role in guiding the work in a team by defining the research project, iteratively "teaching" supervised learning algorithms to recognize appropriate patterns, and interpreting results.

- *Programming.* Staff with strong programming skills and experience managing big data are essential to any machine learning project. Most machine learning projects are programmed in Python or R, though other languages may also be used. Because of the sheer volume of data with which these projects are likely to work, staff in programming and data management roles should have extensive experience in database management and familiarity with ways to streamline a dataset for more efficient processing.

- *IT and Computing Solutions.* IT and cloud computing experts are also required once a machine learning project reaches a certain size. Many machine learning projects will require computing resources far beyond those available in an organization's normal computing environment. On this project, the study anticipated needing one TB of RAM for our prioritization algorithm. Working with such a large volume of data will require most organizations to use cloud computing solutions at an additional cost to the project. To set up these cloud computing solutions, project teams will need to work with specialized IT staff to coordinate with the third-party vendor providing cloud computing tools, build the cloud computing environment within the organization's IT infrastructure, activate servers as necessary, and ensure the cloud computing environment complies with federal data security requirements.

> ***Algorithmic bias and fairness*** refers to the concern that algorithmic decision-making, especially predictive analytics, has the potential to amplify existing disparities and introduce new ones. To be carried out ethically and responsibly, machine learning projects should consider bias and fairness at every stage, from conception through completion. Because machine learning algorithms operate exclusively on the data they are fed, they cannot independently avoid inequitable outcomes if these data reflect systems of bias or inequality. Fortunately, tools exist to monitor and correct algorithmic bias, with data scientists and subject matter experts collaborating to tailor the process to the nuanced context in which a given project exists. Funders of machine learning projects should seek out partners committed to ethical machine learning and the expertise to achieve it. See AlgorithmWatch (https://algorithmwatch.org/en); Awwad et al. (2020); and Mahoney et al. (2020) for more information.

MACHINE LEARNING IS EVOLVING, AND APPROACHES TO THE PRACTICAL CONSIDERATION SURROUNDING THESE METHODS WILL LIKELY EVOLVE AS WELL.

Machine learning is evolving rapidly, and there are many practical considerations machine learning teams will need to address before beginning their work on government contracts, which are often subject to additional regulations and require a greater level of transparency than projects completed in the private sector. As discussed below, much of the legal precedent governing machine learning is changing (see *Authors Guild, Inc. v. Google, Inc..; HiQ Labs, Inc. v. LinkedIn Corporation*) and as the technology that supports machine learning develops, so too will the budgetary considerations surrounding these projects.

Legal and Data Security Considerations

The laws governing the use of machine learning are in various stages of development around the world. In the United States, some aspects of data use and security have clear legal precedent, while others do not. Based on the experience of the CP D&A project, before engaging in any machine learning work, it is important to consider the following questions:

- *Data Access and Use:* Are researchers allowed to access the data they have intended to use for machine learning purposes? Are there any restrictions on that data's use? For example, if scraping third-party websites for information, do their terms of service allow such a scrape? If not, what prohibitions exist? The answers to these questions might depend, in part, on whether these sites can be accessed only after logging in through a password-protected account.
- *Third-Party Vendors:* If using a third-party vendor[8] for any component of the machine learning project, how does the third-party vendor keep data secure? How will data be transferred? Are these methods in line with the data security requirements outlined in our organization's agreement with its funding agency?
- *Data Security:* How will the study keep the data secure? Will it collect and analyze any personally identifiable information (PII)? If not, how will PII be removed from the data set, especially if human researchers cannot

[8] In some cases, a third-party vendor can execute certain aspects of a machine learning project, especially if those aspects require special expertise or enhanced computational power. For this project, the study used SerpAPI, a third-party web-scraping vendor, to conduct the Google searches.

review every data source? If so, does the study have a data security plan and/or Institutional Review Board review to ensure human subjects of research are protected?

- *Funder Requirements:* What requirements or prohibitions around machine learning and data security does the funder have? Does the funder need to approve the use of a third-party vendor for these purposes?

Based on the CP D&A project, the contracting organization's legal team, the third-party vendor (if applicable), and the funder should work in concert to address these questions. If using particular data for the study's intended purposes is limited or prohibited, other data or analysis options may need to be explored.[9] In some cases, a legal precedent that invalidates the prohibitions or limitations listed in a website's terms of service may exist. In others, the third-party vendor may assume any legal liability in the work it conducts for a study.

Budget Considerations

Based on the experience of this project, it is important to accurately budget the cost of a study and know that the nature of these costs may change over time as technology evolves. The study experience indicated machine learning projects have four broad categories of costs: staff time, computing time (if applicable), data access charges (if applicable), and third-party vendor costs (if applicable). Exhibit 4 describes each of these costs.

Exhibit 4: Likely Costs for Machine Learning Projects

Cost Type	Important Considerations	Example
Staff Time	Constitutes the bulk of the costs associated with a machine learning project. Will increase as project and/or task complexity increases, timelines are extended, and/or more experienced staff are required.	Researchers should budget hours for: • Subject matter experts to define key concepts and terms, create research questions and a coding protocol, code data, assess and guide the analysis, and interpret and write up findings • Programming staff to write code and run analyses • IT staff to create and maintain the computing environment • Project leadership to manage the project and communicate with funding agency staff
Computing Time	Costs for cloud computing solutions will be driven by: • The computing power needed (RAM, CPUs) • The amount of data needed to store • Whether an analytic tool is needed and its complexity • The length of time the cloud computing solution is needed The project may incur additional costs if proprietary analysis tools are needed for project completion.	Cloud computing vendors such as Amazon, Google, and Microsoft have complex fee structures that vary based on computing needs and usage. Vendors may charge by the second, minute, or gigabyte. Vendors often have cost calculators that can help project teams gain a rough understanding of the costs they are likely to incur.
Data Access Charges (if applicable)	Some data sources may charge a fee for accessing and using data.	Owners of proprietary data sets may charge a flat or per-record fee for using their data. They may include restrictions around publication in their data use agreements.
Third-Party Vendor Costs	Third-party vendors will charge a fee for their services. Fees may vary by the size or complexity of the work.	SerpAPI charged a tiered monthly fee for scraping Google search results. The "Big Data" plan cost $250 a month and allowed the project team to conduct 30,000 Google searches per month.

[9] As we describe in Section 3, it may make sense for a project to engage in a discovery phase prior to starting machine learning work to assess the data available and the feasibility of the proposed approach.

3. Looking Ahead

The CP D&A experience identifies a number of lessons for future applications of machine learning methods, including those used to answer relevant workforce development policy and research questions. While this exploratory work has been useful, machine learning methods are typically tailored to both the research questions of interest and the data to be used to answer them. The experience of this project suggests that policymakers, government agencies, and other stakeholders interested in using machine learning may first want to execute a "discovery phase" to maximize the efficacy of future machine learning efforts.

Specifically, during a discovery phase, data scientists and subject matter experts could collaborate to generate research questions of interest, explore and assess potential data sources, and create design options that could appropriately apply machine learning techniques to answer questions of interest.

Exhibit 5: Example of Goals and Activities for Discovery Phase for a Workforce-focused Machine Learning Project

Goal	Potential Activities
Identify important workforce development research questions or operational challenges that might be well-answered with machine learning methods.	• Engage in knowledge development activities with workforce development stakeholders to understand their research questions of interest • Convene focus groups with program and data-focused staff to understand: – what information is available, how complete it is, and what it is likely to tell us – any ethical considerations associated with applying machine learning in a given context • Meet with funders and leadership to understand priorities and share findings
Explore data sources internal and (where appropriate) external to DOL and assess their suitability for use in machine learning methods.	• Meet with program and data staff to understand the available quantitative and qualitative data • Obtain a select number of promising internal data sets • If appropriate, explore the feasibility of including external datasets • Assess the suitability of these datasets for a variety of machine learning methods able to answer the research questions of interest
Explore potential legal, budget, or contractual considerations that may impact design decisions.	• Meet with legal representatives to understand funder organization policies on data use and machine learning • Meet with contract and procurement offices to understand potential limitations around contract structure and flexibility around timeline and resources, especially as related to computing resources
Create design options that DOL can pursue in future work.	• Prioritize which questions can be best answered given DOL's interests as well as the data and machine learning methods available • Draft a design memo describing these options • Provide a brief presentation to discuss options with DOL leadership

At the completion of the discovery phase, decisions can be made about the type of machine learning solutions to pursue. The process outlined above is designed to help yield promising opportunities to use machine learning for new research projects and operational improvements. Machine learning methods allow researchers to examine a larger volume and wider breadth of data than could be reasonably reviewed by human-only researchers, and they can help researchers overcome their own biases, preconceptions, and blind spots. When applied appropriately, machine learning has the potential to advance the workforce development field's research and operations goals in ways that were difficult or infeasible before.

Suggested citation: Mills De La Rosa, Siobhan, Nathan Greenstein, Deena Schwartz, and Charlotte Lloyd. (2021). Machine Learning in Workforce Development Research: Lessons and Insights. Rockville, MD: Abt Associates

This report was prepared for the U.S. Department of Labor, Chief Evaluation Office by Abt Associates under Contract Number DOL-1605DC-18-A-0037/1605DC-18-F-00389. The views expressed are those of the authors and should not be attributed to DOL, nor does mention of trade names, commercial products, or organizations imply endorsement of same by the U.S. Government.

References

Agrawal, B., Liu, R., Kokku, R., Chee, Y.-M., Jagmohan, A., Nitta, S., Tan, M., & Sin, S. (2017). 4C: Continuous Cognitive Career Companions. In E. André, R. Baker, X. Hu, Ma. M. T. Rodrigo, & B. du Boulay (Eds.), Artificial Intelligence in Education (pp. 623–629). Springer International Publishing. https://algorithmwatch.org/en/

Alpaydin, E. (2009). *Introduction to Machine Learning*. MIT Press.

Alpaydin, E. (2016). *Machine Learning: The New AI.* MIT Press.

Authors Guild v. Google, Inc., No. 13-4829. United States Court of Appeals for the Second Circuit. 2015. Url: https://law.justia.com/cases/federal/appellate-courts/ca2/13-4829/13-4829-2015-10-16.html

Awwad, Y., Fletcher, R., Frey, D., Gandhi, A., Najafian, M., & Teodorescu, M. (2020). Exploring Fairness in Machine Learning for International Development [Technical Report]. CITE MIT D-Lab. https://dspace.mit.edu/handle/1721.1/126854

Berdanier, C., McComb, C, & Zhu, W. (2020). Natural Language Processing for Theoretical Framework Selection in Engineering Education Research. 2020 IEEE Frontiers in Education Conference (FIE), pp. 1-7, doi: 10.1109/FIE44824.2020.9274115.

Blei, D. M. (2012). Probabilistic Topic Models. *Commun. ACM*, *55*(4), 77–84. https://doi.org/10.1145/2133806.2133826

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

Goldberg, Y. (2017). *Neural Network Methods in Natural Language Processing*. Morgan & Claypool Publishers.

HiQ Labs, Inc. v. LinkedIn Corporation. No. 17-16783. United States Court of Appeals for the Ninth Circuit. 2019. Url: http://cdn.ca9.uscourts.gov/datastore/opinions/2019/09/09/17-16783.pdf

Holzinger A., & Jurisica I. (2014) Knowledge Discovery and Data Mining in Biomedical Informatics: The Future Is in Integrative, Interactive Machine Learning Solutions. In: Holzinger A., Jurisica I. (eds) Interactive Knowledge Discovery and Data Mining in Biomedical Informatics. Lecture Notes in Computer Science, vol 8401. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-662-43968-5_1

Kreimeyer, K., Foster, M., Pandey, A., Arya, N., Halford, G., Jones, S. F., Forshee, R., Walderhaug, M., & Botsis, T. (2017). Natural language processing systems for capturing and standardizing unstructured clinical information: A systematic review. Journal of Biomedical Informatics, 73, 14–29. https://doi.org/10.1016/j.jbi.2017.07.012

Kurdi, M. Z. (2016). *Natural Language Processing and Computational Linguistics: Speech, Morphology and Syntax*. John Wiley & Sons.

Mahoney, T., Varshney, K. R., & Hind, M. (2020). AI Fairness: How to Measure and Reduce Unwanted Bias in Machine Learning. O'Reilly.

Manhattan Strategy Group. (2016). Career Pathways Toolkit: An Enhanced Guide and Workbook for Systems Development. Toolkit prepared by Manhattan Strategy Group. U.S. Department of Labor

Nadkarni, P., Ohno-Machado, L., & Chapman, W. (2011) Natural language processing: an introduction. Journal of the American Medical Informatics Association, 18, 5, 544–551. https://doi.org/10.1136/amiajnl-2011-000464

Sarna, M., & Adam, T. (2020). Evidence on Career Pathways Strategies: Highlights from a Scan of the Research. Career Pathways Brief. Abt Associates. https://www.dol.gov/sites/dolgov/files/OASP/evaluation/pdf/ETA_CareerPathways_Brief_November2020.pdf

Sarna, M., & Strawn, J. (2018). Career Pathways Implementation Synthesis: Career Pathways Design Study. Report prepared by Abt Associates. U.S. Department of Labor. https://www.dol.gov/sites/dolgov/files/OASP/legacy/files/3-Career-Pathways-Implementation-Synthesis.pdf

U.S. Department of Labor. (2016). Career Pathways Toolkit: An Enhanced Guide and Workbook for System Development. https://careerpathways.workforcegps.org/resources/2016/10/20/10/11/Enhanced_Career_Pathways_Toolkit