

# 1

---

## The Dataset Nutrition Label

### *A Framework to Drive Higher Data Quality Standards*

---

SARAH HOLLAND,<sup>1</sup> AHMED HOSNY,<sup>2</sup> SARAH NEWMAN,<sup>3</sup>  
JOSHUA JOSEPH<sup>4</sup> AND KASIA CHMIELINSKI<sup>5</sup>

#### **Abstract**

Data is a fundamental ingredient in building Artificial Intelligence (AI) models and there are direct correlations between data quality and model robustness, fairness and utility. A growing body of research points to AI systems deployed in a wide range of use cases, where algorithms trained on biased, incomplete, or ill-fitting data produce problematic results. Despite the increased critical attention, data interrogation continues to be a challenging task with many issues being difficult to identify and rectify. Algorithms often come under scrutiny only after they are developed and deployed, which exacerbates this problem and underscores the need for better data vetting practices earlier in the development pipeline. We introduce the Dataset Nutrition Label,<sup>6</sup> a diagnostic framework built by the Data Nutrition Project, comprising a label that provides a distilled yet comprehensive overview of dataset ‘ingredients’. The label is designed to be flexible and adaptable; it is comprised of a diverse set of qualitative and quantitative modules generated through multiple statistical and probabilistic modelling backends. Working with the ProPublica dataset ‘Dollars for Docs’, we developed an open source tool<sup>7</sup> consisting of seven sample modules. Consulting such a label prior to AI model development promotes vigorous data interrogation practices, aids

<sup>1</sup> Assembly, MIT Media Lab and Berkman Klein Center at Harvard University.

<sup>2</sup> Dana-Farber Cancer Institute, Harvard Medical Institute.

<sup>3</sup> metaLAB @ Harvard, Berkman Klein Center for Internet & Society, Harvard University.

<sup>4</sup> MIT Quest for Intelligence.

<sup>5</sup> Assembly, MIT Media Lab and Berkman Klein Center at Harvard University.

<sup>6</sup> Available at: <https://datanutrition.media.mit.edu/>.

<sup>7</sup> Available at: <https://ahmedhosny.github.io/datanutrition/>.

in recognising inconsistencies and imbalances, provides an improved means to selecting more appropriate datasets for specific tasks and subsequently increases the overall quality of AI models. We also explore some challenges of the label, including generalising across diverse datasets, as well as discussing research and public policy agendas to further advocate its adoption and ultimately improve the AI development ecosystem.

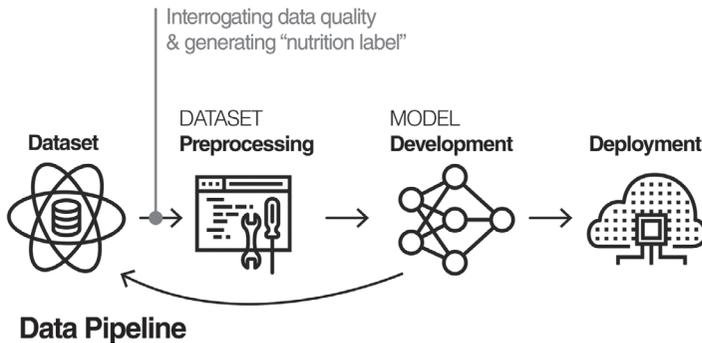
### Keywords

Artificial intelligence, machine learning, data ethics, bias, ethics.

## I. Introduction

Data-driven decision-making systems play an increasingly important role in our lives. These frameworks are built on increasingly sophisticated artificial intelligence (AI) systems and are created and tuned by a growing population of data specialists<sup>8</sup> to arrive at a diversity of decisions: from movie and music recommendations to digital advertisements and mortgage applications.<sup>9</sup> These systems deliver untold societal and economic benefits, but they can also be harmful to individuals and society at large.

**Figure 1.1** Model Development Pipeline



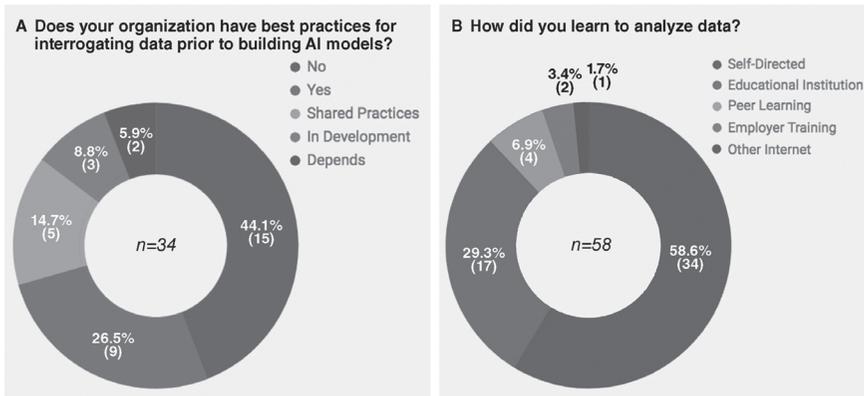
Data is a fundamental ingredient of AI and the quality of a dataset used to build a model will directly influence the outcomes it produces. An AI model trained on problematic data will likely produce problematic outcomes. Examples of

<sup>8</sup>The term 'data specialist' is used instead of 'data scientist' in the interest of using a term that is broadly scoped to include all professionals utilising data in automated decision-making systems: data scientists, analysts, machine learning engineers, model developers, artificial intelligence researchers and a variety of others in this space.

<sup>9</sup>Thomas H Davenport and Jeanne G Harris, 'Automated Decision Making Comes of Age' (2005) 46 *MIT Sloan Management Review* 83.

these include gender bias in language translations surfaced through natural language processing<sup>10</sup> and skin shade bias in facial recognition systems due to non-representative data.<sup>11</sup> Typically, the model development pipeline (Figure 1.1) begins with a question or goal. Within the realm of supervised learning, for instance, a data specialist will curate a labelled dataset of previous answers in response to the guiding question. Such data is then used to train a model to respond in a way that accurately correlates with past occurrences. In this way, past answers are used to forecast the future. This is particularly problematic when outcomes of past events are contaminated with (often unintentional) bias. Models often come under scrutiny only after they are built, trained and deployed. If a model is found to perpetuate a bias – for example, over-indexing for a particular race or gender – the data specialist returns to the development stage to identify and address the issue. This feedback loop is inefficient, costly and does not always mitigate harm; the time and energy of the data specialist is a sunk cost and, if in use, the model deployment may have already produced problematic outcomes. Some of these issues could be avoided by undertaking a thorough interrogation of data at the outset of model development. However, this is still not a widespread practice within AI model development efforts.

**Figure 1.2** (A) Survey results about data analysis best practices in respondents’ organisations and (B) Survey results about how respondents learned to analyse data



We conducted an anonymous online survey (see Figure 1.2), the results of which further lend credence to this problem. Although many (47%) respondents report conducting some form of data analysis prior to model development, most (74%)

<sup>10</sup>Tolga Bolukbasi and others, ‘Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings’, *Advances in Neural Information Processing Systems* (2016).

<sup>11</sup>Joy Buolamwini and Timnit Gebru, ‘Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification’ in Sorelle A Friedler and Christo Wilson (eds), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (PMLR 2018).

indicate that their organisations do not have explicit best practices for such analysis. Fifty-nine per cent of respondents reported relying primarily on experience and self-directed learning (through online tutorials, blogs, academic papers, Stack Overflow and online data competitions) to inform their data analysis methods and practices. This survey indicates that, despite limited current standards, there is widespread interest in improving data analysis practices and making them accessible.

To improve the accuracy and fairness of AI systems, it is imperative that data specialists can assess more quickly the viability and fitness of datasets and more easily find and use better-quality data to train their models. As a proposed solution, we introduce a dataset nutrition label, a diagnostic framework to address and mitigate some of these challenges by providing critical information to data specialists at the point of data analysis. The label thus acts as a first point of contact where decisions regarding the utility and fitness of specific datasets can be made. This is achieved by allowing the recognition of dataset inconsistencies and exclusions as well as promoting dataset interrogation as a crucial and inevitable procedure in the AI model development pipeline – with the ultimate goal of improving the overall quality of AI systems.

We begin with a review of related work, largely drawing from the fields of nutrition and privacy, where labels are a useful mechanism to distill essential information, enable better decision-making and influence best practices. We then discuss the dataset nutrition label prototype, our methodology, demonstration dataset and key results. This is followed by an overview of the benefits of the tool, its potential limitations and ways to mitigate those limitations. We then briefly summarise some future directions, including research and public policy agendas that would further advance the goals of the label. Lastly, we discuss implementation of the prototype and key takeaways.

## II. Labels in Context and Related Work

To inform the development of our prototype and concept, we surveyed the literature for labelling efforts. Labels and warnings are utilised effectively in product safety,<sup>12</sup> pharmaceuticals,<sup>13</sup> energy<sup>14</sup> and material safety.<sup>15</sup> We largely draw from

<sup>12</sup>US Consumer Product Safety Commission. Office of the General Counsel, *Compilation of Statutes Administered by CPSC* (US Consumer Product Safety Commission 1998).

<sup>13</sup>Foster D McClure and United States. *Food and Drug Administration, FDA Nutrition Labeling Manual: A Guide for Developing and Using Databases* (US Food and Drug Administration 1993).

<sup>14</sup>Europäische Union, 'Directive 2009/28/EC of the European Parliament and of the Council of 23 April 2009 on the Promotion of the Use of Energy from Renewable Sources and Amending and Subsequently Repealing Directives 2001/77/EC and 2003/30/EC' (2009) 5 *Official Journal of the European Union* 2009.

<sup>15</sup>Occupational Safety, Health Administration and Others, 'Hazard Communication Standard: Safety Data Sheets' [2012] OSHA Brief.

the fields of nutrition, online privacy and algorithmic accountability as they are particularly salient for our purposes. The former is the canonical example and a long-standing practice subject to significant study while the latter provides valuable insights in the application of a 'nutrition label' in other domains, particularly in subjective contexts and where there is an absence of legal mandates and use is voluntary. Collectively, they elucidate the impacts of labels on audience engagement, education and user decision making.

In 1990, Congress passed the Nutrition Labeling and Education Act (P.L. 101-535), which includes a requirement that certain foodstuffs display a standardised 'Nutrition Facts' label.<sup>16</sup> By mandating the label, vital nutritional facts were communicated in the context of the 'Daily Value' benchmark and consumers could quickly assess nutrition information and more effectively abide by dietary recommendations at the moment of decision.<sup>17</sup> In the nearly three decades since its implementation, several studies have examined the efficacy of the now ubiquitous 'Nutrition Facts' label; these studies include analyses of how consumers use the label<sup>18</sup> and the effect it has had on the market.<sup>19</sup> Though some cast doubt on the benefits of the mandate in light of its cost,<sup>20</sup> most research concludes that the 'Nutrition Facts' label has had a positive impact.<sup>21</sup> Surveys demonstrate widespread consumer awareness of the label and its influence in decision making around food, despite a relatively short time since the passage of the Nutrition Labeling and Education Act.<sup>22</sup> According to the International Food Information Council, more than 80 per cent of consumers reported they looked at the 'Nutrition Facts' label when deciding what foods to purchase or consume and only 4 per cent reported never using the label.<sup>23</sup> Five years after the mandate, the Food Marketing Institute found that about one-third of consumers stopped buying a food because of what they read on the label.<sup>24</sup> With regard to the information contained on the

<sup>16</sup> United States. Congress. House. Committee on Energy and Commerce, *Nutrition Labeling and Education Act of 1990: Report (to Accompany H.R. 3562) (including Cost Estimate of the Congressional Budget Office)* (US Government Printing Office 1990).

<sup>17</sup> Ibid; Siva K Balasubramanian and Catherine Cole, 'Consumers' Search and Use of Nutrition Information: The Challenge and Promise of the Nutrition Labeling and Education Act' (2002) 66 *Journal of Marketing* 112; Joanne F Guthrie and others, 'Who Uses Nutrition Labeling and What Effects Does Label Use Have on Diet Quality?' (1995) 27 *Journal of Nutrition Education* 163.

<sup>18</sup> Balasubramanian and Cole (n 17); Guthrie and others (n 17).

<sup>19</sup> Bruce A Silverglade, 'The Nutrition Labeling and Education Act: Progress to Date and Challenges for the Future' (1996) 15 *Journal of Public Policy & Marketing* 148.

<sup>20</sup> Paul J Petrucci, 'Consumer and Marketing Implications of Information Provision: The Case of the Nutrition Labeling and Education Act of 1990' (1996) 15 *Journal of Public Policy & Marketing* 150.

<sup>21</sup> Mario F Teisl, Alan S Levy and Others, 'Does Nutrition Labeling Lead to Healthier Eating?' (1997) 28 *Journal of Food Distribution Research* 18; Andreas C Drichoutis, Panagiotis Lazaridis and Rodolfo M Nayga Jr, 'Consumers' Use of Nutritional Labels: A Review of Research Studies and Issues' (2006) *Academy of Marketing Science Review* 1.

<sup>22</sup> United States Congress. House Committee on Energy and Commerce (n 16).

<sup>23</sup> Susan Borra, 'Consumer Perspectives on Food Labels' (2006) 83 *The American Journal of Clinical Nutrition* 1235S.

<sup>24</sup> United States Congress. House Committee on Energy and Commerce (n 16).

label and consumer understanding, researchers found that ‘label format and inclusion of (external) reference value information appear to have (positive) effects on consumer perceptions and evaluations’,<sup>25</sup> but consumers indicated confusion about the ‘Daily Value’ comparison, suggesting that more information about the source and reliability of ground truth information would be useful.<sup>26</sup> The literature focuses primarily on the impact to consumers rather than on industry operations such as production and advertising. However, the significant impact of reported sales and marketing materials on consumers<sup>27</sup> provides a foundation for further inquiry into how this has affected the greater food industry.

In the field of privacy and privacy disclosures, the nutrition label serves as a useful point of reference and inspiration.<sup>28</sup> Researchers at Carnegie Mellon and Microsoft created the ‘Privacy Nutrition Label’ to better surface essential privacy information to assist consumer decision making with regard to the collection, use and sharing of personal information.<sup>29</sup> The ‘Privacy Nutrition Label’ operates much like ‘Nutrition Facts’ and sits atop existing disclosures. It improves the functionality of the Platform for Privacy Notices, a machine-readable format developed by the World Wide Web Consortium, itself an effort to standardise and improve the legibility of privacy policies.<sup>30</sup> User surveys that tested the ‘Privacy Nutrition Label’ against alternative formats found that the label outperformed alternatives with ‘significant positive effects on the accuracy and speed of information finding and reader enjoyment with privacy policies’ as well as improved consumer understanding.<sup>31</sup>

Ranking and scoring algorithms also pose challenges in terms of their complexity, opacity and sensitivity to the influence of data. End users and even model developers face difficulty in interpreting an algorithm and its ranking outputs and this difficulty is further compounded when the model and the data on which it is trained is proprietary or otherwise confidential, as is often the case. ‘Ranking Facts’ is a web-based system that generates a ‘nutrition label’ for scoring and ranking algorithms based on factors or ‘widgets’ to communicate an algorithm’s methodology or output.<sup>32</sup> Here, the label serves more as an interpretability tool

<sup>25</sup> Scot Burton, Abhijit Biswas and Richard Netemeyer, ‘Effects of Alternative Nutrition Label Formats and Nutrition Reference Information on Consumer Perceptions, Comprehension and Product Evaluations’ (1994) 13 *Journal of Public Policy & Marketing* 36.

<sup>26</sup> Borra (n 22).

<sup>27</sup> Silverglade (n 19).

<sup>28</sup> Corey A Ciochetti, ‘The Future of Privacy Policies: A Privacy Nutrition Label Filled with Fair Information Practices’ (2008) 26 *The John Marshall Journal of Computer & Information Law* 1.

<sup>29</sup> Patrick Gage Kelley and others, ‘A Nutrition Label for Privacy’, *Proceedings of the 5th Symposium on Usable Privacy and Security* (Association for Computer Machinery (ACM) 2009).

<sup>30</sup> Patrick Gage Kelley and others, ‘Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach’, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (ACM2010).

<sup>31</sup> Kelley and others (n 29); Kelley and others (n 30).

<sup>32</sup> Ke Yang and others, ‘A Nutritional Label for Rankings’. Available at: [www.cs.drexel.edu/~julia/documents/cr-sigmod-demo-2018.pdf](http://www.cs.drexel.edu/~julia/documents/cr-sigmod-demo-2018.pdf).

than as a summary of information as the 'Nutrition Facts' and 'Privacy Nutrition Label' operate. The widgets work together, not modularly, to assess the algorithm on author-created categories of transparency, fairness, stability and diversity. The demonstration scenarios for using real datasets from college rankings, criminal risk assessment and financial services establish that the label is potentially applicable to a diverse range of domains. This lends credence to the potential utility in other fields as well, including the rapidly evolving field of AI.

More recently, in an effort to improve transparency, accountability and outcomes of AI systems, AI researchers have proposed methods for standardising practices and communicating information about the data itself.

The first draws from computer hardware and industry safety standards where datasheets are an industry-wide standard. In datasets, however, they are a novel concept. Datasheets are functionally comparable to the label concept and, like labels that by and large objectively surface empirical information, can often include other information such as recommended uses which are more subjective. 'Datasheets for Datasets', a proposal from researchers at Microsoft Research, Georgia Tech, University of Maryland and the AI Now Institute, seeks to standardise information about public datasets, commercial APIs and pretrained models. The proposed datasheet includes dataset provenance, key characteristics, relevant regulations and test results, but also significant yet more subjective information such as potential bias, strengths and weaknesses of the dataset, API, or model and suggested uses.<sup>33</sup> As domain experts, dataset, API and model creators would be responsible for creating the datasheets, not end users or other parties.

We are also aware of a forthcoming study from the field of natural language processing (NLP), 'Data Statements for NLP: Toward Mitigating System Bias and Enabling Better Science'.<sup>34</sup> The researchers seek to address ethics, exclusion and bias issues in NLP systems. Borrowing from similar practices in other fields of practice, the position paper puts forward the concept and practice of 'data statements' which are qualitative summaries that provide detailed information and important context about the populations the datasets represent. The information contained in data statements can be used to surface potential mismatches between the populations used to train a system and the populations in planned use prior to deployment, to help diagnose sources of bias that are discovered in deployed systems and to help understand how experimental results might generalise. The authors suggest that data statements should eventually become required practice for system documentation and academic publications for NLP systems and should be extended to other data types (eg image data) albeit with tailored schema.

We take a different, yet complementary, approach. We hypothesise that the concept of a 'nutrition label' for datasets is an effective means to provide a scalable

<sup>33</sup>Timnit Gebru and others, 'Datasheets for Datasets' (2018). Available at: <http://arxiv.org/abs/1803.09010>.

<sup>34</sup>n/a, 'Data Statements for NLP: Toward Mitigating System Bias and Enabling Better Science'.

and efficient tool to improve the process of dataset interrogation and analysis prior to and during model development. In supporting our hypothesis, we created a prototype, a dataset nutrition label. Three goals drive this work. First, to inform and improve data specialists' selection and interrogation of datasets and to prompt critical analysis. Consequently, data specialists are the primary intended audience. Second, to gain traction as a practical, readily deployable tool, we prioritise efficiency and flexibility. To that end, we do not suggest one specific approach to the label or charge one specific community with creating the label. Rather, our prototype is modular and the underlying framework is one that anyone can utilise. Lastly, we leverage probabilistic computing tools to surface potential corollaries, anomalies and proxies. This is particularly beneficial because resolving these issues requires excess development time and can lead to undesired correlations in trained models.

### III. Methods

Some assumptions are made to focus our prototyping efforts. Only tabular data is considered. Additionally, we limit our explorations to datasets <10k rows. This allows for a narrower scope and deeper analysis. The label's first contribution lies in the standard format it provides for metadata communication. This works to address weaknesses in the most common format for tabular data exchange: comma separated values, or the '.csv' format. Despite its simple plain-text nature, portability and interoperability,<sup>35</sup> the lack of additional .csv metadata describing how data should be interpreted, validated and displayed is, perhaps, its biggest drawback. As early as 2015, the World Wide Web Consortium published recommendations on 'Metadata Vocabulary for Tabular Data' and 'Access methods for CSV Metadata'.<sup>36</sup> However, the adoption of these recommendations within the data science community is not widespread. The label also builds on existing data science practices: directly following the acquisition of a dataset, most data specialists often enter an 'exploratory phase'. This can be seen, for instance, on web-hosted machine learning competition platforms such as Kaggle and involves understanding dataset distributions through histograms and other basic statistics. The label attempts to provide these statistics 'out of the box' with the hopes of shortening model development lead times. The label also aims to provide insights from advanced probabilistic modelling backends for more advanced users. While targeted mainly at a professional audience, many of the modules do not require expert training for interpretation and can thus be utilised in a public-facing version of the label.

<sup>35</sup> Andrew Hunt and David Thomas, *The Pragmatic Programmer: From Journeyman to Master* (Boston: Addison-Wesley Professional 2000).

<sup>36</sup> 'Metadata Vocabulary for Tabular Data'. Available at: [www.w3.org/TR/tabular-metadata/](http://www.w3.org/TR/tabular-metadata/) (last accessed 17 April 2018); 'Model for Tabular Data and Metadata on the Web'. Available at: [www.w3.org/TR/tabular-data-model/](http://www.w3.org/TR/tabular-data-model/) (last accessed 17 April 2018).

**Table 1.1** Seven modules of the label, together with their description, role, contents and attributes, highlighting the potential diversity of information included in a label

Module Name	Description	Contents	Module Characteristic – Level Required				
			Technical Expertise	Manual Effort	Subjectivity	Interactivity	Data Exposure
<b>Metadata</b>	Meta information. This module is the only required module. It represents the absolute minimum information to be presented	Filename, file format, URL, domain, keywords, type, dataset size, % of missing cells, licence, release date, collection range, description	Low	High	Low	Low	Low
<b>Provenance</b>	Information regarding the origin and lineage of the dataset	Source and author contact information with version history	Low	High	Low	Low	Low
<b>Variables</b>	Descriptions of each variable (column) in the dataset	Textual descriptions	Low	High	Medium	Low	Medium
<b>Statistics</b>	Simple statistics for all variables, in addition to stratifications into ordinal, nominal, continuous, and discrete	Least/most frequent entries, min/max, median, mean, etc	Medium	Low	Low	Low	Medium
<b>Pair Plots</b>	Distributions and linear correlations between 2 chosen variables	Histograms and heatmaps	Medium	Low	Low	High	High
<b>Probabilistic Model</b>	Synthetic data generated using distribution hypotheses from which the data was drawn – leverages a probabilistic programming backend	Histograms and other statistical plots	High	Medium	High	Low	High
<b>Ground Truth Correlations</b>	Linear correlations between a chosen variable in the dataset and variables from other datasets considered to be “ground truth”, such as Census Data	Heatmaps	Medium	Medium	Low	Low	High

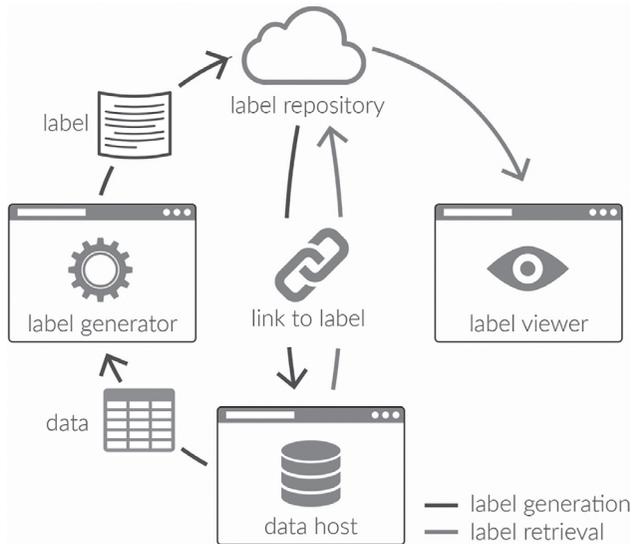
## A. Modular Architecture

The label is designed in an extensible fashion with multiple distinct components that we refer to as ‘modules’ (see Table 1.1). The modules are stand-alone, allowing for greater flexibility as arrangements of different modules can be used for different types of datasets. This format also caters to a wide range of requirements and information available for a specific dataset. During label generation and subsequent updates, it also accommodates data specialists of different backgrounds and technical skill levels.

Modules (see Table 1.1) range from the purely non-technical, such as the Metadata module, to the highly technical, such as the Probabilistic Computing module. Some modules require manual effort to generate, such as those that provide qualitative descriptions of the data (Metadata, Provenance, Variables), while others can ideally be the result of an automated process (Statistics, Pair Plots). Modules also vary in their subjectivity, especially where there exists a reliance on the label author to identify which questions should be asked of the data and in what way (eg Probabilistic Computing). Many of the example modules are also interactive, highlighting a crucial benefit of a label living on a platform (such as a web page) that supports user interaction. This allows label users to interrogate various dataset aspects with great flexibility and free of preconceived notions developed during label generation. Lastly, some modules could be designed to act as proxies for their corresponding dataset as they do not expose the underlying data. This could be key when dealing with proprietary datasets, as much of this data will not or cannot be released to the public based on intellectual property or other constraints. Other modules expose information such as distribution metrics which, in theory, would allow adversaries to approximate the dataset contents. The choice of module(s) is thus based on the availability of information, level of willingness and effort volunteered to document the dataset and privacy concerns.

The list of modules currently examined in this study, while not exhaustive, provides a solid representation of the kinds of flexibility supported by the label framework. Other modules considered for future iterations or additional datasets include but are not limited to: a comments section for users to interact with authors of the label for feedback or other purposes; an extension of the Provenance section that includes the versioning history and change logs of the dataset and associated labels over time, similar to Git; a privacy-focused module that indicates any sensitive information and whether the data was collected with consent; and finally, a usage tracking module that documents data utilisation and references using some form of identifier, similar to the Digital Object Identifier<sup>37</sup> and associated citation systems in scientific publishing.

<sup>37</sup> Rajesh Chandrakar, ‘Digital Object Identifier System: An Overview’ (2006) 24 *The Electronic Library* 445.

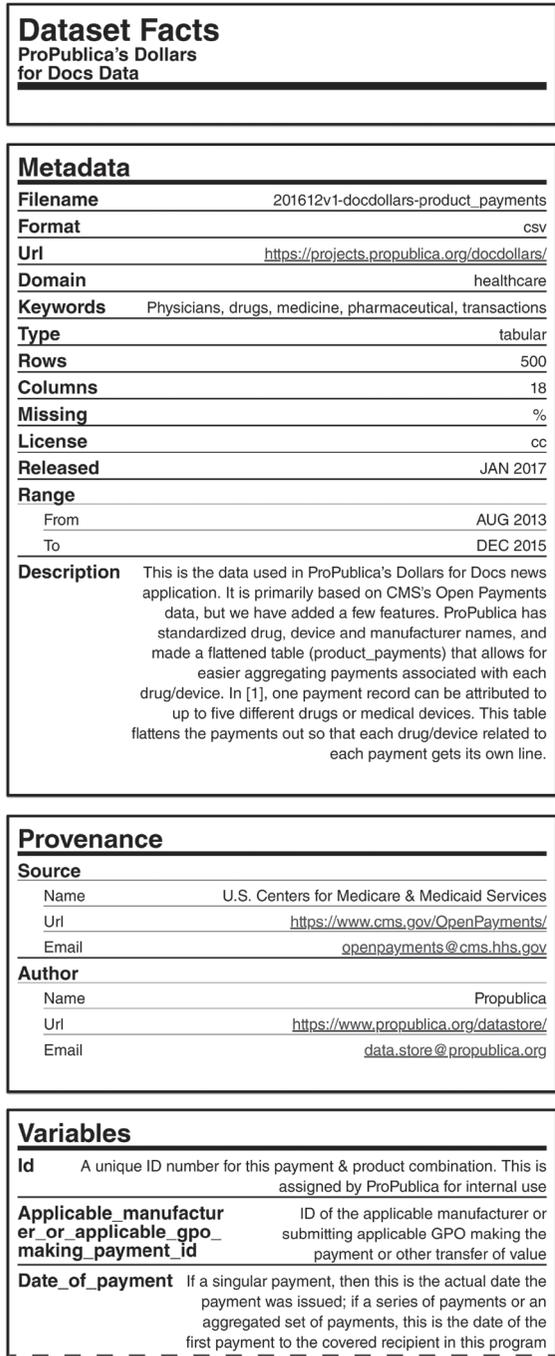
**Figure 1.3** Architecture of the proposed nutrition label ecosystem

## B. Web-based Application

The label is envisioned as a digital object that can be both generated and viewed by web-based applications. The label ecosystem comprises two main components: a label maker and a label viewer (see Figure 1.3). Given a specific dataset, the label maker application allows users to select the desired modules and generate them. While the generation of some modules is fully automated, some require human input (see Table 1.1). For instance, the Metadata module mainly requires explicit input, while the Pair Plots module can be generated automatically from the dataset. The label generator pre-populates as many fields as possible and alerts users to those requiring action. The label itself lives in a .json format, as one that is human readable and well supported. The label can then be viewed within the label viewer application where formatting is carried out to achieve the desired user interface and user interaction effects. In terms of visual appearance and design, format and typeface requirements of the ‘Nutrition Facts’ label<sup>38</sup> is used. These guidelines, such as the all black font color on white contrasting background, are optimised for clarity and conciseness. Design changes are anticipated in further iterations and should be informed by user testing.

<sup>38</sup> US Food, Drug Administration and Others, ‘A Food Labeling Guide’ (Center for Food Safety & Applied Nutrition, 1999).

**Figure 1.4** Prototype label demonstrating the metadata, provenance and variables modules



## C. Backends

Simple statistical analyses involving the generation of histograms, distribution information and linear correlations are carried out directly in the browser, given tabular datasets of <100K rows. Server-side processing is thus reserved for more specialised and sophisticated analyses requiring additional computational power. Such processing could run multiple backends with the ultimate aim of providing the label authors with a diverse set of options, fueled by the plethora of tools developed by research groups for automating the generation of summaries, insights and understandings of datasets. The label thus becomes a medium for the continuous deployment and testing of these tools. A somewhat recent and particularly powerful example of this is probabilistic computing and specifically, BayesDB,<sup>39</sup> an open source platform developed by researchers at MIT. With minimal modeling and programming effort, BayesDB enables inference of a model that captures the structure underlying the data and generates statistical summaries based on such structure.

## IV. Results

To test the concept generally and the modular framework specifically, we built a prototype with a dataset that included information about people and was maintained by an organisation invested in better understanding the data. This combination of factors provides necessary information and access to build a wide variety of modules, including those that require full knowledge of the data and the ability to contact the organisation that maintains the dataset. We were granted access to the ‘Dollars for Docs’ database from ProPublica, an independent, nonprofit newsroom that produces investigative journalism in the public interest.<sup>40</sup> The dataset, which contains payments to doctors and teaching hospitals from pharmaceutical and medical device companies over a two-year time period (August 2013–December 2015), was originally released by the US Centers for Medicare and Medicaid Services (CMS) and compiled by ProPublica into a single, comprehensive database.

The resulting prototype successfully demonstrates how disparate modules can be built on a specific dataset in order to highlight multiple, complementary facets of the data, ideally to be leveraged for further investigation by data specialists through the use of additional tools and strategies. The prototype label includes seven modules (see Table 1.1). The Metadata, Provenance and Variables modules

<sup>39</sup> Vikash Mansinghka and others, ‘BayesDB: A Probabilistic Programming System for Querying the Probable Implications of Data’ (2015). Available at: <http://arxiv.org/abs/1512.05006>.

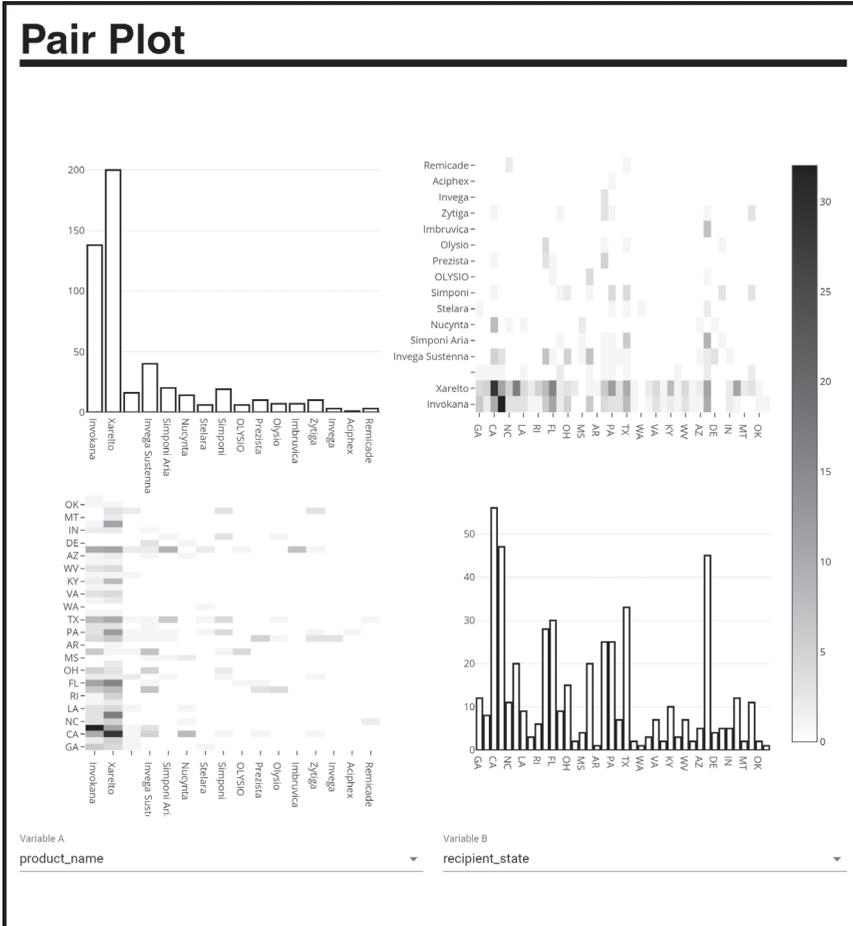
<sup>40</sup> Available at: <https://projects.propublica.org/docdollars/>.

**Figure 1.5** Prototype label demonstrating the Statistics module, splitting the variables into four groups: ordinal, nominal, continuous and discrete

Statistics									
<b>Ordinal</b>									
name	type	count	uniqueEntries	mostFrequent	leastFrequent	missing			
id	number	500	488 including mi...	missing value (13)	multiple detected	2.60%			
applicable_man...	number	500	4	100000000232 (...)	multiple detected	0%			
date_of_payment	date	500	213 including mi...	missing value (27)	multiple detected	5.40%			
general_transac...	number	500	467 including mi...	missing value (34)	multiple detected	6.80%			
program_year	number	500	2 including miss...	2014 (495)	missing value (5)	1.00%			
<b>Nominal</b>									
name	type	count	uniqueEntries	mostFrequent	leastFrequent	missing			
product_name	string	500	16 including mis...	Xarelto (200)	Aciphex (1)	3.20%			
original_product...	string	500	15	Xarelto (212)	Aciphex (1)	0%			
product_ndc	number	500	21 including mis...	5045857810 (201)	multiple detected	5.00%			
product_is_drug	boolean	500	2 including miss...	t (492)	missing value (8)	1.60%			
payment_has_...	boolean	500	3 including miss...	f (267)	missing value (29)	5.80%			
teaching_hospit...	number	500	2 including miss...	0 (484)	missing value (36)	7.20%			
physician_profil...	number	500	230 including mi...	missing value (32)	multiple detected	6.40%			
recipient_state	string	500	40	CA (56)	multiple detected	0%			
applicable_man...	string	500	5 including miss...	Janssen Pharm...	multiple detected	7.00%			
teaching_hospit...	number	500	2 including miss...	0 (481)	missing value (19)	3.80%			
product_slug	string	500	15 including mis...	drug-xarelto (196)	drug-aciphex (1)	8.20%			
<b>Continuous</b>									
name	type	count	min	median	max	mean	standardD...	missing	zeros
total_amo...	number	500	0.14	14.00	5000	134.21	501.99	9.40%	0%
<b>Discrete</b>									
name	type	count	min	median	max	mean	standardD...	missing	zeros
number_o...	number	500	1	1.00	1	1.00	0.00	4.80%	0%

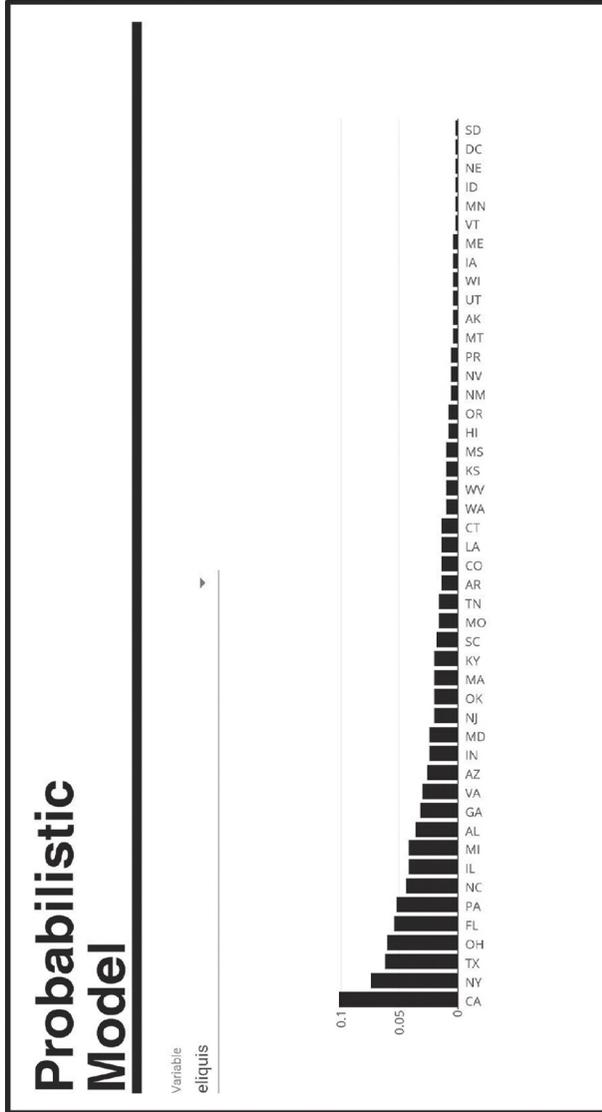
(see Figure 1.4) provide as-is dataset information. They mirror information submitted by the label authors as well as provide a standard format for both the generation and consumption of such data. The Statistics module (see Figure 1.5) starts to offer a glimpse into the dataset distributions. For instance, the skewness of a 500-row dataset subset towards a particular drug ‘Xarelto’ can be quickly identified as the most frequent entry under the variable ‘product name’ and ‘Aciphex’ as the least frequent entry. The Pair Plot module (see Figure 1.6) starts to introduce

**Figure 1.6** Prototype label demonstrating the Pair Plot module and highlighting the interactive dropdown menus for selecting variables



interactivity into the label where the viewer is able to choose the variable pair being compared to one another. A specialist building a model predicting marketing spend in each state, for example, may choose to compare 'recipient\_state' and 'total\_amount\_of\_payment\_USdollars' and will observe that some states (CA, NY) are more highly correlated with spend. In this case, the specialist would probably normalise for population as the next step beyond consulting the label in order to identify anomalous spending trends.

Figure 1.7 Prototype label demonstrating the Probabilistic Model module and showcasing a hypothetical distribution for payments made towards the drug 'Eliquis' across different states

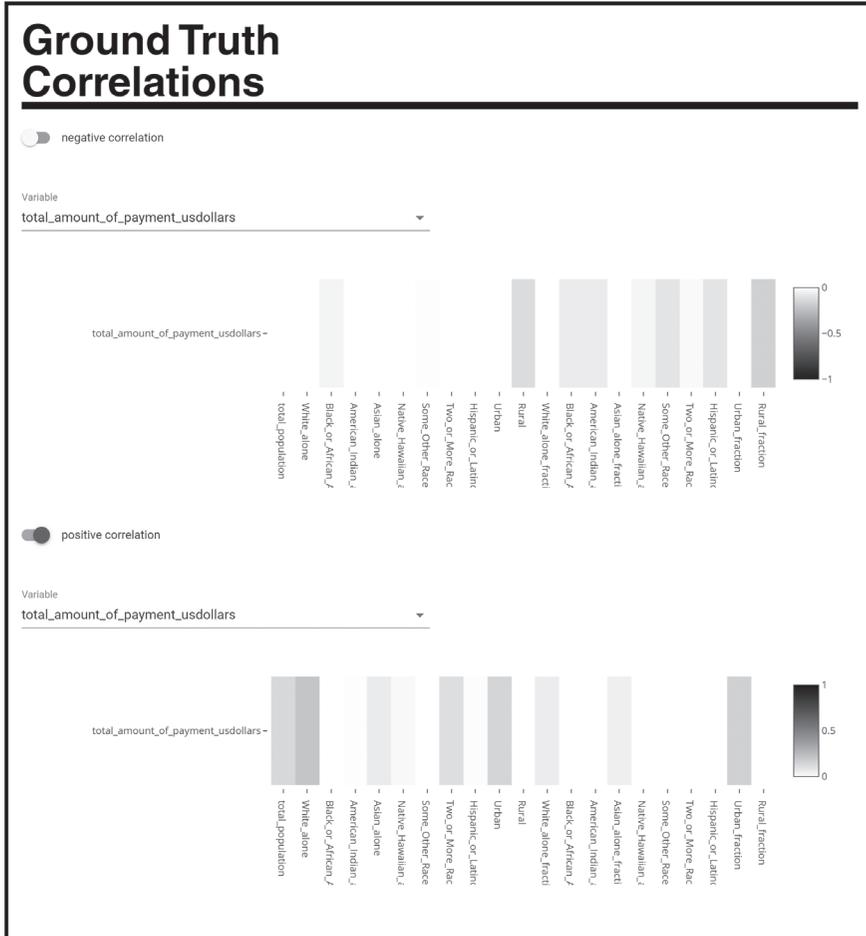


While all modules thus far investigate the dataset itself, the Probabilistic Model module (see Figure 1.7) attempts to generate synthetic data by utilising the aforementioned BayesDB backend. Computed from an inferred generative model, this module allows for the full benefits of Bayesian analysis,<sup>41</sup> such as interpretability of inferences, coping with missing data and robustness to outliers and regions of sparse data. In this specific use case, an underrepresented drug is chosen from the dataset and the probability of this drug receiving a payment in different states is inferred. With the inevitable variation in the representation of different groups in datasets, such analyses are of great utility in extracting insights – even from relatively small sample sizes. A quick toggle indicates that the top few states for marketing spend are likely the same few states – with a few exceptions, including that New Jersey is likely to receive a lot more money for marketing activities relating to the drug Xarelto. Again, this information only acts as a flag for the ‘what’; specialists will ideally continue to investigate the data in order to identify the ‘why’.

It is unavoidable that datasets collected from the real world have relationships to demographics that the data specialist or other entities do not wish to propagate into the learned model and the inferences produced from it. For example, is a variable or an aggregate of a variable strongly correlated with the Hispanic population in a given region? To surface relationships like this, it is often necessary to explicitly compute a comparison between the dataset and demographic ‘ground truth’ data, which is a task that can be both time consuming and challenging. The Ground Truth Correlation module (see Figure 1.8) provides the data specialist initial evidence as to whether such relationships are likely, thus warranting further analysis. In order to surface any anomalies in the demographic distribution of these variables, we selected the 2010 US Census data as ‘ground truth’ for zip code and race. The module then correlates zip code Census data with the dataset and calculates the Pearson correlation between demographics and field aggregates. To demonstrate its utility, the label (see Figure 1.8 top) highlights the negative correlations between the (sum of the) amount of payment field and demographics. A second example (see Figure 1.8 bottom), highlights the positive correlation between a ‘*spend\_per\_person*’ aggregate and demographics. This module demonstrates, in a straightforward way, specific anomalous relationships in the data that the data specialist should pay attention to during model training. In the prototype, we observe a slight positive correlation between white zip codes and payments and a slight negative correlation between rural zip codes and payments. Toggling to per person spend underscores similar overall trends.

<sup>41</sup> James O Berger, *Statistical Decision Theory and Bayesian Analysis* (Springer Science & Business Media 2013).

**Figure 1.8** The negative (top) and positive (bottom) correlations to demographics produced by the Ground Truth Correlations module



## V. Discussion

The label offers many benefits. Overall, it prompts critical questions and interrogation in the preprocessing phase of model development. It also expedites decision making, which saves time in the overall model development phase without sacrificing the quality or thoroughness of the data interrogation itself, perhaps encouraging better practices at scale. These benefits apply across the spectrum of data specialists' skill and experience but are particularly useful for those new to the field or less attuned to concerns around bias and algorithmic accountability. First, the label creates a pre-generated 'floor' for basic data interrogation in the data selection phase. It also indicates key dataset attributes in a standardised format.

This gives data specialists a distilled yet comprehensive overview of the ‘ingredients’ of the dataset, which allows for a quick and effective comparison of multiple datasets before committing to one for further investigation. It also enables the data specialist to better understand and ascertain the fitness of a dataset by scanning missing values, summary statistics of the data, correlations or proxies and other important factors. As a result, the data specialist may discard a problematic dataset or work to improve its viability prior to utilising it.

Improved dataset selection affords a secondary benefit: higher quality models. The label provides data specialists improved means by which to interrogate the selected dataset during model development, previously a costly and onerous enterprise. The Ground Truth Correlation module, in particular, provides a helpful point of reference for the data specialist before model completion and surfaces issues such as surprising variable correlations, missing data, anomalous data distributions, or other factors that could reinforce or perpetuate bias in the dataset. Addressing these factors in the model creation and training phase saves cost, time and effort and could also prevent bad outcomes early on, rather than addressing them after the fact.

The label is built with scalability in mind and with an eye towards standardisation. The modular framework provides flexibility for dataset authors and publishers to identify the ‘right’ kind and amount of information to include in a label; over time, this could become a set of domain-specific best practices. The interactivity of the label also permits flexibility, as insights about the dataset may arise over time. For example, the ground truth data used for comparison could evolve, rendering a previously unsuitable dataset suitable. Interactive labels also give data specialists the ability to dive further into anomalous data, rather than simply accepting static information provided by the label author. With some modules more subjective in nature and with a range of domain expertise across data specialists, this is particularly important. For advanced data specialists, the flexible label backend makes it easy to ‘plug-in’ more complex engines. Such complex backends can provide different statistical tools; for example, the Probabilistic Computing module makes it possible to investigate low frequency variables by generating synthetic data. Synthetic data gives data specialists the ability to address incomplete data and opens a potential path to privacy preserving data usage.<sup>42</sup>

Lastly, the label functioning as a proxy for the dataset itself is an intriguing, even if distant, possibility. Increased calls for accountability of AI systems demand investigation of datasets used in training models, but disclosing those datasets, even to a limited audience, may pose risks to privacy, security and intellectual property that calls this approach into question. If the label is able to convey the

<sup>42</sup> Mohan HS Surendra H, ‘A Review of Synthetic Data Generation Methods For Privacy Preserving Data Publishing’ (2017) 6 *International Journal of Scientific & Technology Research* 95–101. Available at: [www.ijstr.org/final-print/mar2017/A-Review-Of-Synthetic-Data-Generation-Methods-For-Privacy-Preserving-Data-Publishing.pdf](http://www.ijstr.org/final-print/mar2017/A-Review-Of-Synthetic-Data-Generation-Methods-For-Privacy-Preserving-Data-Publishing.pdf).

information essential for accountability in the dataset without disclosing the data itself, it would provide a valuable and much-needed auditing tool<sup>43</sup> for AI systems while still preserving privacy, security and proprietary information.

## A. Limitations and Mitigations

There are challenges to our approach. The extensive variety of datasets used to build models raises important questions around whether the label can generalise across data and dataset type, size, composition and in different domains and furthermore, whether a data specialist or domain expert will need to be involved in the creation of a label across these different datasets. This could arise in an instance where important semantic information is atypically labeled and would be challenging to interpret automatically, such as if the field for zip code in a dataset had a custom field for 'geographic area'. A data specialist or domain expert may also be required when building a label for sensitive or proprietary data, which may be accessible only to those who built the dataset and not accessible to the public. Building the label as a modular system somewhat mitigates the complication of requiring input from a domain expert, as the framework can adapt to domain-specific best practices and can easily support the generation of different types of labels based on access. Within the Provenance module, it may be necessary or helpful to surface who made the label and what relationship they have to the dataset.

The veracity and usefulness of the Ground Truth Comparison module depends on the accuracy of the 'ground Truth' dataset, which serves as a benchmark standard and is considered objective, accurate and provable and with clear provenance. However, problematic ground truth data may lead to futile or even harmful comparisons. Without a realistic way to eliminate bias in all datasets, a mitigating step is to build labels for ground truth datasets themselves. If these labels include community feedback and comment modules, dataset authors can address the issues directly.

Further investigation is necessary to understand the feasibility and desirability of using the label as a proxy for proprietary datasets. This would likely require that the dataset creator or controller create the label. Another challenge is that the label might not prompt the right questions or reactions for the data specialist, leaving certain biased data undetected. Analyses of machine bias indicate that zip codes often proxy for race, but many others proxies still exist, especially as the models themselves approach levels of complexity that are difficult or impossible for humans to comprehend and new or unexpected proxies emerge. Integrating new methods or tools to help identify proxies will be important to the industry

<sup>43</sup> Select Committee on Artificial Intelligence, 'AI in the UK: Ready, Willing and Able?' Available at: <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>.

and our hope is that the label will be flexible in such a way that these tools can be leveraged to create additional modules as they become available.

Finally, design of the label itself will require additional attention to determine the appropriate amount of information for presentation, comprehension and adoption. As made clear in the work on Privacy Nutrition Label,<sup>44</sup> design is a key element in the efficacy of the label. It is worth investigating and testing the most effective presentation to drive adoption.

## VI. Future Directions

This paper and prototype are the first step toward bringing together a wide range of data specialists, from those creating and publishing datasets to those utilising datasets to build models, in order to improve the quality of datasets used in AI-generated models.

Deeper research and iteration will be necessary as we continue to build additional prototypes of the label. Creating a ‘nutrition label’ for datasets is nascent and requires additional investigations about what information (in the form of modules or otherwise) is useful and practical to include. Based on the relatively small reach of our survey, we also recommend that a more rigorous survey be conducted to more accurately identify needs, as the survey we administered was limited in its reach and disproportionately indexed to American and European respondents working in the private sector. The information pertinent to a data specialist will also shift based on the domain of the data, necessitating the building of additional prototypes for different kinds of datasets. The opportunities afforded by complex machine learning tools such as BayesDB in the creation of additional modules deserve more fulsome exploration to maximise the usability and usefulness of the label.

Through building relationships with dataset publishers and circulating the label, we hope to identify not only additional datasets for prototypes, but also to launch our label on open datasets so that we can study the impact of the label on the use of and conversation around the data. We will consider collaborations with colleagues from industry and academia to further drive this work, building knowledge around the impediments to adoption and considering ways that regulatory frameworks could further support the creation of a best practice or standard.

In terms of the label ecosystem, the existence of a label for any given dataset could be notated using a mark or symbol, such as the ‘Conformité Européenne’ (CE) mark used by the European Union,<sup>45</sup> on the author’s or dataset host’s webpage. Clicking on the mark would then navigate to the label viewer application and fetch the corresponding label from a central repository where all labels are hosted.

<sup>44</sup> Kelley and others (n 29).

<sup>45</sup> David Hanson, *CE Marking, Product Standards and World Trade* (Edward Elgar Publishing 2005).

Such a centralised archive of labels would allow for generating usage statistics, least and most used modules and eventually help inform future label iterations. More importantly, a repository of this sort could act as an index of datasets without hosting the datasets themselves. For instance, API calls to such a repository could help locate datasets with queries like ‘MIT license dataset for facial recognition with >100k samples’.

Beyond its utility as a tool, the label could also drive a change in norms. Through using the label, data specialists will build a habit around questioning datasets through analysis and interrogation techniques, even if a particular dataset does not include a label. In time, the label will facilitate an environment that encourages a broad spectrum of dataset creators, cleaners, publishers and users to create labels to publish alongside their datasets. This would lead to better identification of issues with data and bias, or inappropriate data collection practices, which in turn would increase data and dataset quality overall.

Looking beyond the label itself, there are longer term opportunities for this framework and the data science community. Decisions made around the authorship and ownership model for the label will be critical to the overall direction of the project; who will create these labels going forward and who will maintain them? Will there be a single place where all labels live or from where they are all linked? Additional future directions could include: building a public consortium or governing body to consider standards across the industry; creating curriculum for those collecting and working with datasets; and further exploration of appropriate ground truth data.

## VII. Conclusions

In an effort to improve the current state of practice of data analysis, we created the dataset nutrition label, a diagnostic framework that provides a concise yet robust and standardised view of the core components of a dataset. We use the ProPublica Dollars for Docs dataset to create the label prototype.

The label serves as a proof of concept for several conceptual questions, beginning with the general feasibility of an extensible and diverse modular framework. It also confirms the possibility of mixing qualitative and quantitative modules that leverage different statistical and probabilistic modelling backend technologies in the same overall user experience. The label integrates both static and interactive modules, underscoring the importance of using an interactive platform (such as a website) for the distribution of the label itself. Together, this promises flexibility, scalability and adaptability.

With the label, data specialists can efficiently compare, select and interrogate datasets. Additionally, certain modules afford the ability to check for issues with the dataset before and during model development, surface anomalies and potentially dangerous proxies and find new insights into the data at hand. As a

result, data specialists have a better, more efficient process of data interrogation, which will produce higher quality AI models. The label is a useful, practical, timely and necessary intervention in the development of AI models and a first step in a broader effort toward improving the outcomes of AI systems that play an increasingly central role in our lives.

## Acknowledgements

We are grateful to the ProPublica team, including Celeste LeCompte, Ryann Jones, Scott Klein and Hannah Fresques, for their generosity in providing the Dollars for Docs dataset and for their assistance throughout prototype development. We are also grateful to the BayesDB team in the Probabilistic Computing Group at MIT, including Vikash Mansinghka, Sara Rendtorff-Smith and Ulrich Schaechtle for their valuable time, assistance and ongoing advice with regard to integrating the BayesDB backend with the label. Thanks also go to Patrick Gage Kelley for bringing key work to our attention and for his constructive feedback, as well as the 2018 Assembly Cohort and Advisory Board, in particular Matt Taylor, Jack Clark, Rachel Kalmar, Kathy Pham, James Mickens, Andy Ellis and Nathan Freitas; the City of Boston Office of New Urban Mechanics; and Eric Breck and Mahima Pushkarna of Google Brain for productive and insightful discussions. This work was made possible by the Assembly programme led by Jonathan Zittrain of the Berkman Klein Center for Internet & Society and Joi Ito of the MIT Media Lab.

## References

- Balasubramanian SK and Cole C, 'Consumers' Search and Use of Nutrition Information: The Challenge and Promise of the Nutrition Labeling and Education Act' (2002) 66 *Journal of Marketing* 112.
- Berger JO, *Statistical Decision Theory and Bayesian Analysis* (Springer Science & Business Media 2013).
- Bolukbasi T and others, 'Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings', *Advances in Neural Information Processing Systems* (2016).
- Borra S, 'Consumer Perspectives on Food Labels' (2006) 83 *The American Journal of Clinical Nutrition* 1235S.
- Buolamwini J and Gebu T, 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification' in Sorelle A Friedler and Christo Wilson (eds), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (PMLR 2018).
- Burton S, Biswas A and Netemeyer R, 'Effects of Alternative Nutrition Label Formats and Nutrition Reference Information on Consumer Perceptions, Comprehension and Product Evaluations' (1994) 13 *Journal of Public Policy & Marketing* 36.

- Chandrakar R, 'Digital Object Identifier System: An Overview' (2006) 24 *The Electronic Library* 445.
- Ciocchetti CA, 'The Future of Privacy Policies: A Privacy Nutrition Label Filled with Fair Information Practices' (2008) 26 *The John Marshall Journal of Computer and Information Law* 1.
- Davenport TH and Harris JG, 'Automated Decision Making Comes of Age' (2005) 46 *MIT Sloan Management Review* 83.
- Drichoutis AC, Lazaridis P and Nayga RM Jr, 'Consumers' Use of Nutritional Labels: A Review of Research Studies and Issues' (2006) 2006 *Academy of Marketing Science Review* 1.
- Geburu T and others, 'Datasheets for Datasets' (2018). Available at: <http://arxiv.org/abs/1803.09010>.
- Guthrie JF and others, 'Who Uses Nutrition Labeling and What Effects Does Label Use Have on Diet Quality?' (1995) 27 *Journal of Nutrition Education* 163.
- Hanson D, CE Marking, Product Standards and World Trade (Edward Elgar Publishing 2005).
- Hunt A and Thomas D, *The Pragmatic Programmer: From Journeyman to Master* (Addison-Wesley Professional 2000).
- Kelley PG and others, 'A Nutrition Label for Privacy', Proceedings of the 5th Symposium on Usable Privacy and Security (Association for Computer Machinery 2009).
- , 'Standardising Privacy Notices: An Online Study of the Nutrition Label Approach', Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Association for Computer Machinery 2010)
- Mansinghka V and others, 'Bayes DB: A Probabilistic Programming System for Querying the Probable Implications of Data' (2015). Available at: <http://arxiv.org/abs/1512.05006>.
- McClure FD and United States Food and Drug Administration, 'FDA Nutrition Labeling Manual: A Guide for Developing and Using Databases' (US Food and Drug Administration 1993).
- 'Metadata Vocabulary for Tabular Data'. Available at: [www.w3.org/TR/tabular-metadata/](http://www.w3.org/TR/tabular-metadata/) (last accessed 17 April 2018).
- 'Model for Tabular Data and Metadata on the Web'. Available at: [www.w3.org/TR/tabular-data-model/](http://www.w3.org/TR/tabular-data-model/) (last accessed 17 April 2018).
- n/a, 'Data Statements for NLP: Toward Mitigating System Bias and Enabling Better Science'
- Petrucelli PJ, 'Consumer and Marketing Implications of Information Provision: The Case of the Nutrition Labeling and Education Act of 1990' (1996) 15 *Journal of Public Policy & Marketing* 150.
- Safety O, Administration H and Others, 'Hazard Communication Standard: Safety Data Sheets' [2012] OSHA Brief.
- Select Committee on Artificial Intelligence, 'AI in the UK: Ready, Willing and Able?' (House of Lords). Available at: <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>.
- Silverglade BA, 'The Nutrition Labeling and Education Act: Progress to Date and Challenges for the Future' (1996) 15 *Journal of Public Policy & Marketing* 148.
- Surendra H MHS, 'A Review of Synthetic Data Generation Methods for Privacy Preserving Data Publishing' (2017) 6(3) *International Journal of Scientific and Technology Research* 95–101. Available at: [www.ijstr.org/final-print/mar2017/A-Review-Of-Synthetic-Data-Generation-Methods-For-Privacy-Preserving-Data-Publishing.pdf](http://www.ijstr.org/final-print/mar2017/A-Review-Of-Synthetic-Data-Generation-Methods-For-Privacy-Preserving-Data-Publishing.pdf).

- Teisl MF, Levy AS and Others, 'Does Nutrition Labeling Lead to Healthier Eating?' (1997) 28 *Journal of Food Distribution Research* 18.
- Union E, 'Directive 2009/28/EC of the European Parliament and of the Council of 23 April 2009 on the Promotion of the Use of Energy from Renewable Sources and Amending and Subsequently Repealing Directives 2001/77/EC and 2003/30/EC' (2009) 5 *Official Journal of the European Union* 2009.
- United States. Congress. House. Committee on Energy and Commerce, Nutrition Labeling and Education Act of 1990: Report (to Accompany HR 3562) (including Cost Estimate of the Congressional Budget Office) (Washington DC: US Government Printing Office 1990).
- US Consumer Product Safety Commission. Office of the General Counsel, Compilation of Statutes Administered by CPSC (Bethesda, MD: US Consumer Product Safety Commission 1998).
- US Food Administration and Others, 'A Food Labeling Guide' [1999] Center for Food Safety and Applied Nutrition.
- Yang K and others, 'A Nutritional Label for Rankings.' Available at: [www.cs.drexel.edu/~julia/documents/cr-sigmod-demo-2018.pdf](http://www.cs.drexel.edu/~julia/documents/cr-sigmod-demo-2018.pdf).

