

Responsible Data Management for Human Resources

Dimitrios Vogiatzis**

The American College of Greece, Deree
& NCSR "Demokritos"
Athens, Greece
dimitrv@acg.edu

Olivia Kyriakidou

The American College of Greece, Deree
Athens, Greece
OKyriakidou@acg.edu

ABSTRACT

The human resources (HR) departments rely increasingly on recommender systems (RS) for most of their processes, such as recruiting, selecting and developing their employees. However, the RS often discriminate unfairly based on biases in data that may perpetuate and enhance existing biases and in the work place. An important part of an HR department is a the data ecosystem, comprising raw and derived data, related to potentially different stakeholders while being subject to laws, and regulations. In this work we propose the characteristics of a data ecosystem that will facilitate data transparency through traceability as a way of detecting potential biases in the data.

CCS CONCEPTS

• **Information systems** → **Data management systems**; • **Social and professional topics** → **Employment issues**; *User characteristics*.

KEYWORDS

Human Resources, Fairness, Bias, Data Ecosystems

1 INTRODUCTION

Recommender systems (RS) are widely used by Human Resources (HR) departments to facilitate their business processes from the point of time optimization, but also from the perspective of minimizing human intervention, in an effort of achieving fairness. RS can be applied in the recruitment, hiring, promotion of employees, etc. For instance they can be used in matching CVs against job posts, in ranking CVs which determines the order of interviews or even in comparison of CVs against past CVs of employees that are deemed as successful. A RS can also be used for segmenting job applications into categories, detecting and recording long terms trends etc. Moreover, they can be used by prospective employees that seek employment.

Although RS seem to remove human intervention by automating HR processes, often but not exclusively, through advanced machine learning algorithms, still segments of the population can be discriminated against. The data upon which the analysis is based on may contain biases towards age groups, gender, ethnic origin etc. The data stem from specific data items, specific data features, data distributions, data sampling methods. The bias in data can be also very subtle and difficult to detect as it may appear in derived data stemming from an analytics process.

*Both authors contributed equally to this research.

Eventually a RS in HR is a information system that is directly related to the professional life of people, and as such it should be subject to ethical and legal regulations, apart from the technical ones, like prediction accuracy. ACM¹ and IEEE² have issued codes of ethics that refer to the need for fairness in Information Systems. In particular section 1.4 of ACM code of ethics is entitled *Be fair and take action not to discriminate* and section II of the IEEE code of ethics states: *To treat all persons fairly and with respect, to not engage in harassment or discrimination, and to avoid injuring others.*

In reality RS are often fraught with elements of discrimination and unfairness. The unfairness may stem from biases in the data that may misrepresent the actual population and subsequent analytic algorithm often amplify the data biases. Lack of fairness can have potential legal consequences, especially in employment as it might violate anti-discrimination laws. Also it might have financial consequences as usage of such systems might drop. See also [10] for a recent tutorial on the origin and form of fairness in RS.

A data ecosystem is a network of data in potentially many forms (e.g. unstructured, structured) as well as accompanying rules that permit their acquisition, storage, maintenance, and retrieval. The data ecosystems is of potential interest to many stakeholders, including data providers, and data users that try to create value out the ecosystems. A data ecosystem includes metadata, as well as legal, organizational or ethical regulations. Moreover, the ecosystems evolve as their constituent components change. Finally, derived data also form part of the ecosystems. For instance clusters, predictions etc. are examples of derived data that are produced by statistical or machine learning methods. See [15] for an overview of data ecosystems.

Our contribution is to focus on the data component of a RS and examine how a data ecosystem would facilitate data transparency through data traceability so that potential biases are made explicit or a least easier to track and detect. Our approach is based on a similar work for data transparency in the biomedical domain [11].

2 RELATED WORK

Responsible data management has been discussed in the context of automated decision systems (ADS), which are systems that make decisions about humans that might affect their socio-economic life [17], [18]. They authors refer to the ethical challenges faced in all phases of a data science pipeline and the need for a fair, transparent and responsible data management.

Our current work focuses on the data representation part of an ADS, which is essentially an RS. In particular, we refer to the features of a data ecosystem and how it can be supported with

¹<https://www.acm.org/code-of-ethics>

²<https://www.ieee.org/about/corporate/governance/p7-8.html>

semantic web technologies, and their relevance to the issues of fairness in HR.

A very similar approach that we propose in the current work has been developed for a biomedical system in the context of the EU funded project BigMedilytics³ for a lung-cancer pilot application. The pilot integrates structured and unstructured information, open and sensitive data in a knowledge graph. This constitutes an example of a data ecosystem.

3 MOTIVATING EXAMPLES

Next we mention some examples that indicate the form of bias in raw data, in data associations, as well as in derived data produced by machine learning algorithms. The examples refer to HR department cases.

Biased Data based on human behavioral biases. HR algorithmic recommendations may sustain existing inequities when they are trained on data that do not include specific groups of individuals [9]. For example, many selection algorithms try to identify the criteria that characterize the ideal employee and use them for the selection of newcomers. For this task they utilize performance data that identify the best performing employees within the organization and then identify the traits that distinguish them. However, there is the danger that if the performance data favor men due to existing biases within the organization [16], then the selection algorithm might include gender as a preferred characteristic for the ideal candidate and prefer men rather than women applicants. In this sense, existing biases could be reified by limiting the number of certain groups and possibly underrepresented groups who are alerted, selected, and hired for specific job openings [6]. Moreover, HR recommendation systems utilized for the automated screening of candidates' CVs against certain preferred selection criteria may also generate biased results when they are trained on data from past hiring decisions that are based on individual, organizational and structural biases against certain underrepresented groups of employees [14]. [4]. The use of natural language processing (NLP) tools in chatbots that evaluate candidates' competencies and fit to the job and the organization may also preserve existing societal inequities when they are trained on biased data and exclude certain categories of candidates. The association of African-American names with negative feelings and female names with the household and non-technical jobs has been already documented in the literature [19].

Proxies. Recommendation systems can replicate biases in other subtle ways, especially through the use of proxies. Certain hiring criteria could serve as proxies for categorizing individuals in specific groups and drive discrimination. For example, the use of gaps in employment as a hiring criterion could discriminate against women applicants as women disproportionately leave the workplace to provide child or elderly care [1]. Moreover, job matching platforms and job recommendation systems use proxies for "relevance" that reproduce biases. Such systems, for example, could show to women specific jobs at specific hierarchical levels (e.g., senior or junior positions in management) according to their own search history but also according to the search history of women similar to them. Accordingly, they might end up with fewer recommendations for

senior positions if themselves and others look tend to look for lower-level jobs [4]. Proxies are also included in HR data that train employee selection recommendation systems in order to offer the most appropriate remuneration package to prospective employees. Such suggestions however may reinforce gender or racial pay gaps especially when they reflect the existence of strong proxies that signal for certain gender representations (e.g., male employees as breadwinners) and status inequalities.

Facial analysis that is used in virtual interviews may also create disparate impact on specific sub-groups of employees across gender and racial lines. In [5] it was shown that the faces of women with darker skin cannot be reliably recognized by facial analysis systems as well as the emotions of people with disabilities and in different cultural contexts [3]. Finally, in employee selection, most recruiters use a number of candidate characteristics as proxies of culture fit [7], defined as the degree to which the values of the individual match those of the organization. However, there is the danger that these proxies will become hard rules ignoring their subjective character and in this way exclude certain individuals who are thought apriori that they do not "fit" the organizational culture.

Segregation of individuals. Biases could also persist when algorithms segregate employees into groups drawing inferences about individuals from their group memberships. Selection recommendation systems, for instance, may erroneously attribute to people with disabilities [20] certain characteristics based on their group membership without properly assessing the candidates and consequently offer lower status job positions. Moreover, categorizing individuals into certain gender groups could unfairly marginalize non-binary and transgender employees while their classification into certain race groups could signify status inequalities [12].

Human computer interaction. Most HR recommendation systems run on platforms that require employees' and candidates' active involvement with them which is determined merely by the rules set by the platform that control all processes [13]. For instance, job candidates do not have any control over how their application will be presented to possible employers and they have to provide all the required information by the platform if they want to be considered for future job opportunities. Moreover, employee selection recommendation systems tend to present numerical rankings of candidates to employers generating the perception that there are actual substantial differences between the candidates for a certain position, while in reality the differences might be minimal [2].

4 REQUIREMENTS FOR A DATA ECOSYSTEM IN HR

Next, based on the previously mentioned examples we sketch the requirements that would be necessary for a data ecosystem in HR.

Data management requirements: The data ecosystem should allow data sharing for structured (e.g. CSV files) and unstructured data (e.g. text). The data should be accessible and retrievable by all stakeholders. Also the data has to be of high quality. For instance, data items that have missing values could be rejected or data items that are very old. As an example we could mention a CV that that does not contain any information about education or past employment. The data

³<https://www.bigmedilytics.eu/>

management requirements fall into the following categories: **DM1: Data management of multiple document types** should be supported. **DM2: Quality of data items** should be supported at all levels of data pipeline, e.g. at the raw data, but also at derived data.

Organizational requirements: The data should be stored, accessed and processed according to the organization's rules, and regulations. The organization requirements fall into the following categories: **O1: Data governance** should be enforced by the organization. Thus the data acquisition process, the data storage, and retention, access rights, and data obsolescence are items related to data governance. The HR department may have business rules that stipulate the recruitment policy, and what will be the requested documents. **O2: Data sovereignty** which specifies who owns the original data, the derived data, and to what purpose. This will increase the trust in the system. For instance, it will be clearer of how a submitted CV will be handled.

Legal & Ethical requirements: The data management should be in accordance with the requirements of the European GDPR.⁴ Moreover, the data management should address bias. For instance, the execution of algorithms should be independent of sensitive attributes (like ethnicity, age, gender). In addition, the data should be owned and used for the intended purposes. For instance, CVs of job applicants should not be used to generate business value by selling them without their consent. Finally, traceability is an important aspect of the data that essentially allows to know where the data were obtained from, and how they were obtained.

The above can be summarized into the following ethical requirements: **E1: Data protection & ownership** which specifies the extent of ownership for each stakeholder.

E2: Sensitive attributes which clearly states the sensitive attributes, with the foresight that they should be used by prediction algorithms. Typically, they represent age, gender, ethnic background etc. The sensitive attributes are typically associated with the provisions of GDPR.

E3: discrimination attributes they may lead to discrimination in non-obvious ways. For instance the name of a job applicant might inadvertently facilitate discrimination as it may reveal ethnic origin. Moreover, some derived attributes fall in this category, for instance employment gaps.

5 DESIGN OF A DATA ECOSYSTEM

We present in detail the concept of a data ecosystem, that will serve as the infrastructure for an HR. A data ecosystem (DE) can be defined as a 4-tuple: $DE = \langle \text{Data Sets}, \text{Data Operators}, \text{Meta-Data}, \text{Mappings} \rangle$ [8].

Data sets: the ecosystem is composed of potentially multiple data sets. Data sets can be comprised of structured or unstructured information; also, they have different formats, e.g., CSV, JSON or tabular relations, and can be managed using different management systems.

Data operators: the set of operators that can be executed against the data sets. For instance, anonymization, data quality checks, recency checks can be considered as data operators.

Meta-Data: provide the semantics of the data stored in the data sets of the data ecosystem. It comprises:

- (1) *A Domain ontology*, which provides a unified view of the concepts, relationships, and constraints of the domain of knowledge. It associates formal elements from the domain ontology to concepts. For instance, a specific job post and a specific applicant can be part of the concepts in a domain ontology.
- (2) *Properties* enable the definition of data quality, provenance, and data access regulations of the data in the ecosystem. For instance, last updated and other non-domain properties (quality etc).
- (3) *Descriptions* of the main characteristics of a data set. No specific formal language or vocabulary is required; in fact, a data set could be described using natural language. For instance, Data set D is a collection of CVs and cover letters.

Mappings expressing correspondences among the different components of a data ecosystem. The mappings are as follows:

Mappings between ontologies: they represent associations between the concepts in the different ontologies that compose the domain ontology of the ecosystem. For instance, if there can be a mapping between the personnel ontology, and the candidate employees ontology.

Mappings between data sets: they represent relations among data sets of the ecosystem and the domain ontology.

6 DATA ECOSYSTEM IN HR

The role of a data ecosystem (DE) is to provide an explicit description of the data and applicable operations on them through meta-data and mapping rules. Next we provide some examples that refer to the usage of the elements of the DE as referred to in the previous section. The description of a data ecosystem that we describe next does not cover all the activities HR, but rather it addresses some essential parts that refer to recruitment and hiring of employees. Thus we will assume a scenario where there are applicants CVs and job posts. The DE for this example can be depicted in Figure 1.

One of the data sources are the CVs of the applicants. Typically, they contain textual information, possibly with some keywords (e.g. education, past employment) which can be helpful as annotations. Thus a CV represents a piece of unstructured or partially structured information.

A *Data Operator* can implement an NLP process to extract structured information from the CV. Typically named entity recognition (NER) and relation extraction (RE) will have to be performed, resulting in triplets comprising two entities and a relation. The named entities (NE) in CVs can be things like: skills, past employment, educational achievements and demographic data. The relations connect the NE to the person in question, while being labeled with time annotations. This will form the job applicant's graph. The extracted entities will then be annotated with *meta data* derived from a *domain ontology*, commonly described in OWL. For example

⁴<https://gdpr-info.eu/>

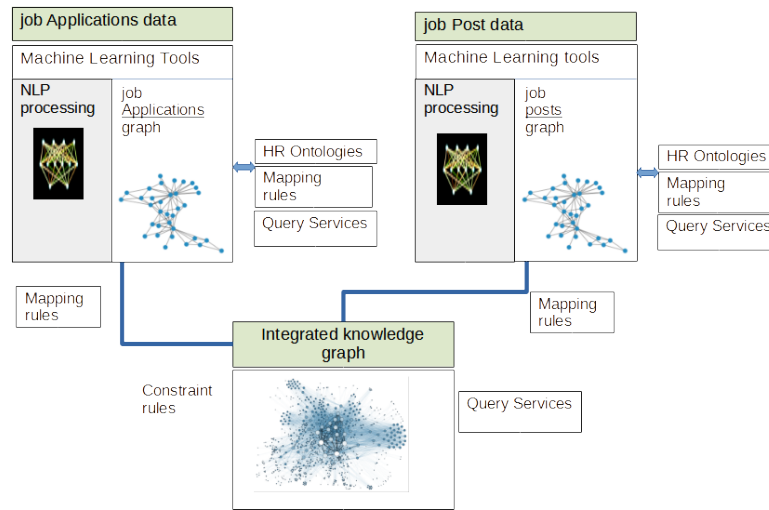


Figure 1: The HR Data Ecosystem

the NAICS⁵ can be used to characterise the entities that refer to the industry of employment.

The NER and RE processes can on occasion be of low precision. The *Meta-data properties* can represent the quality of the NLP process as a numerical score per NE or per relation.

To the best of our knowledge there is not single ontology that is complete enough to annotate a CV for the requirements of an HR. For instance, it may be necessary apart from NAICS to use also resumeRDF,⁶ and the Human Resources Ontology.⁷ This results in the need also to have mappings between the ontologies for the common concepts (i.e. classes) and for the common object properties. The mapping rules can be stated in RML.⁸

The second major source of information is the job post, which is typically in textual form, possibly split in sections each with a meaningful keyword (like company culture, required skills etc.). This usually constitutes a partially structured piece of information. Likewise with the case of CVs information has to be extracted in the forms of triplets, resulting in the job posts graph. However, it may be not necessary to extract structure from all parts of the document. For instance a company's culture could fall under the *Meta-Data Descriptions* data set.

Finally the merging of the two graphs in the integrated knowledge graph can also be achieved with mapping rules. The mapping rules, as well all the ontology selection, and possibly expansion to be designed in cooperation of a knowledge engineer with a representative of HR department.

The issue of detecting possible *bias* in the data can be assisted through data transparency, especially at the stage of NE annotation. Thus CV attributes can be split into *sensitive* and *non-sensitive* ones. The former comprising name, gender, ethnic origin, age etc. whereas the latter would comprise entities like education, or skills.

The distinction between attributes can be represented for instance as classes by expanding one of the existing ontologies. Thus it will be clearer what attributes should be used from subsequent machine learning algorithms that perform job recommendations.

Subsequently a similar distinction can be made between *soft* and *hard skills* in job posts. Naturally, this will require *Data operators* to split the skills into two classes. This will facilitate an association of soft and hard skills to the level of seniority of the position, and to the applicants' gender. This can reveal subtle forms of biases.

Finally, *business regulations* and regulations derived from *ethical data management* can be set as constraints on the integrate knowledge graph, and be expressed in the SHACL⁹ language.

Typically the DE can accessed via SPARQL endpoints. Normally, the end user has access via web services accessible through a dashboard. The web services can also allow for different user roles, thus implementing data access control.

7 CONCLUSIONS

In the current work we proposed a framework for responsible data management for a human resources department. The framework is based on the concept of a DE, that comprises data, meta-data and data operators. It can be implemented with semantic technologies (RDF Schema, OWL, RML rules, etc.). The implementation of a data ecosystem will require a substantial investment both from a knowledge engineering and the HR perspective. The benefits can be important, especially in the field of data transparency. Moreover, a DE can also facilitate the deployment of explainable machine learning algorithms.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Deree - The American College of Greece in the current article.

⁵North American Industry Classification System (NAICS) <https://www.census.gov/naics/>

⁶<http://rdfs.org/resume-rdf/>

⁷<https://github.com/motapinto/cv-ontology/blob/main/cv-ontology.owl>

⁸<https://rml.io/specs/rml/>

⁹<https://www.w3.org/TR/shacl/>

REFERENCES

- [1] Ifeoma Ajunwa. 2020. The Paradox of Automation as Anti-Bias Intervention, 41 Cardozo, L.
- [2] Ifeoma Ajunwa and Daniel Greene. 2019. Platforms at work: Automated hiring platforms and other new intermediaries in the organization of work. In *Work and labor in the digital age*. Emerald Publishing Limited.
- [3] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. 2019. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20, 1 (2019), 1–68.
- [4] Miranda Bogen and Aaron Rieke. 2018. Help wanted: An examination of hiring algorithms, equity, and bias. (2018).
- [5] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [6] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Balanced neighborhoods for multi-sided fairness in recommendation. In *Conference on Fairness, Accountability and Transparency*. PMLR, 202–214.
- [7] Hege H Bye, Henrik Herrebrøden, Gunnhild J Hjetland, Guro Ø Røyset, and Linda L Westby. 2014. Stereotypes of Norwegian social groups. *Scandinavian Journal of Psychology* 55, 5 (2014), 469–476.
- [8] Cinzia Capiello, Avigdor Gal, Matthias Jarke, and Jakob Rehof. 2020. Data ecosystems: sovereign data exchange among organizations (Dagstuhl Seminar 19391). In *Dagstuhl Reports*, Vol. 9. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [9] Kate Crawford. 2013. The hidden biases in big data. *Harvard business review* 1, 4 (2013).
- [10] Michael D Ekstrand, Robin Burke, and Fernando Diaz. 2019. Fairness and discrimination in recommendation and retrieval. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 576–577.
- [11] Sandra Geisler, Maria-Esther Vidal, Cinzia Capiello, Bernadette Farias Loscio, Avigdor Gal, Matthias Jarke, Maurizio Lenzerini, Paolo Missier, Boris Otto, Elda Paja, et al. 2021. Knowledge-driven Data Ecosystems Towards Data Transparency. *arXiv preprint arXiv:2105.09312* (2021).
- [12] Os Keyes. 2018. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–22.
- [13] Karen Levy and Solon Barocas. 2017. Designing against discrimination in online markets. *Berkeley Technology Law Journal* 32, 3 (2017), 1183–1238.
- [14] Kirsten Martin. 2019. Ethical implications and accountability of algorithms. *Journal of Business Ethics* 160, 4 (2019), 835–850.
- [15] Marcelo Iury S Oliveira and Bernadette Farias Lóscio. 2018. What is a data ecosystem?. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age*. 1–9.
- [16] Lauren A Rivera. 2015. Go with your gut: Emotion and evaluation in job interviews. *American journal of sociology* 120, 5 (2015), 1339–1389.
- [17] Julia Stoyanovich, Bill Howe, Serge Abiteboul, Gerome Miklau, Arnaud Sahuguet, and Gerhard Weikum. 2017. Fides: Towards a platform for responsible data science. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. 1–6.
- [18] Julia Stoyanovich, Bill Howe, and HV Jagadish. 2020. Responsible data management. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3474–3488.
- [19] Adam Sutton, Thomas Lansdall-Welfare, and Nello Cristianini. 2018. Biased embeddings from wild data: Measuring, understanding and removing. In *International Symposium on Intelligent Data Analysis*. Springer, 328–339.
- [20] Shari Trewin. 2018. AI fairness for people with disabilities: Point of view. *arXiv preprint arXiv:1811.10670* (2018).