ELECTRIC POWER
RESEARCH INSTITUTE

# A PRIMER ON

## DATA GOVERNANCE FOR A RESPONSIBLE ARTIFICIAL INTELLIGENCE IN THE POWER INDUSTRY

Learn More: www.ai.epri.com 800.313.3774

Sign up to receive updates about AI.EPRI!

January 2020

AI.EPRI

# Data Governance for a Responsible AI in the Power Industry

## Summary

Today's electrical grid is going through a massive shift driven by decarbonization, decentralization, and digitization. Digitization, in particular, has resulted in new and significant data streams from different parts of the grid. These data streams provide an opportunity for deployment of analytics tools to benefit the grid further. As one of these powerful tools, artificial intelligence (AI) has tremendous potential to generate actionable intelligence from the data streams.

With the new data streams come unique challenges and concerns, particularly when sharing and using data for developing AI tools. These challenges and concerns are not specific to the power industry. In many industries, such as finance and healthcare, new guidelines are being developed to address how data are shared and used for different applications and algorithms as well as what models are being developed based on these data streams. Moreover, policymakers are working to establish appropriate instructions for using AI "responsibly" in various applications.

In this primer, we define what "responsible AI" means, and we summarize the efforts to define guidelines and polices around it. Then, we focus on some best practices and considerations for one component of responsible AI—data governance. However, it should be noted that the paper presents the points of view of different subject matter experts in this industry, and the paper is not meant to advocate any standard or policy.

## Introduction

The electric power industry is going through rapid and massive changes to adapt itself to society's and customers' expectations, which are being reflected in decarbonization, decentralization, and digitization. To address these expectations, the power industry has been modernizing itself for decades. Digitization in generating stations and installment of smart meters are just two examples of such efforts. Moreover, the rise of the Internet of Things and connected devices has enabled customers to be more involved in the integrated grid and has made them more aware of their role in shaping the future of the grid.

An immediate outcome of all this modernizing is the generation of a large volume of data streams that must be handled carefully to produce meaningful insights that can benefit the grid and customers. AI is believed to have tremendous capabilities [1] to help the industry use the created data to make better decisions and to benefit the grid and customers in a fast and cost-efficient manner.

Other industries have already explored and implemented AI. However, many utilities in the power industry are skeptical of how AI can be used to its full benefit without jeopardizing the security and safety of the grid or violating customers' privacy. Moreover, utility customers are now more aware of the value of their data, and they demand higher standards when it comes to using their data.

**Responsible AI:** the practice of creating machine learning algorithms that are governable, sustainable, and accountable (see Figure 1).

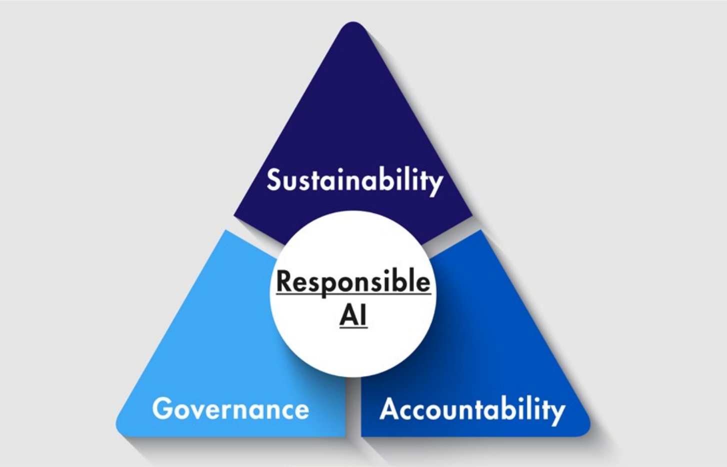

*Figure 1. The components of responsible AI*

## Responsible AI

This section defines responsible AI and describes its different components. Different terms are used to refer to responsible AI, including ethical AI, which might vary in scope and definition. For instance, transparency and explainability [3] of AI are popular terms in discussions about the ethics of AI (but outside the scope of this primer).

# Definition of Responsible AI

Given the broad application of AI, there are many definitions of responsible AI, depending on the industry. However, it can be briefly defined as the practice of creating machine learning algorithms that are governable, sustainable, and accountable. In other words, "they do what they are supposed to do, remain reliable over time and are well controlled and auditable" [4].

There are many different angles to each of these elements. For instance, considerations such as a human-centric design, explainability, and continuous streams of training data fall into one or more of these categories. In this primer, our main concern is data governance because the other components of responsible AI are in their infancy in the power and many other industries.

## 1. Accountability

AI models are expected to automate many processes and decrease or eliminate the involvement of humans. However, it is not yet clear who should be accountable for the potential consequences if these processes do not function as intended. For instance, let's consider a smart thermostat. If, for any reason, the thermostat does not function as it is trained to and causes harm to the household residents or a spike in an electric bill, who or what should be held responsible? The most obvious answer is the manufacturer. However, as AI algorithms become more intelligent, the manufacturer can argue that data provided by the residents have led to the situation, thereby making the customers accountable.

Even though the legal and policy aspects of accountable AI have yet to be developed, following practices—for instance, making sure that the data for models come from unbiased and compliant sources—could move AI models toward being more responsible. Avoiding biases in the models and moving toward explainable AI are other practices. This will help us to understand why an AI model makes an undesirable decision. Moreover, one needs to make sure that objectivity in interpreting the results is maintained. This requires a collaborative approach to model development, involving a human as a supervisor in different steps of model development as well as ensuring that multiple checks and balances are in place regarding the behavior of the AI model.

## 2. Sustainability

Sustainable AI refers to the practice of building a framework in which secure and continuous data access is guaranteed. Moreover, for sustainable AI, this framework should be able to continuously monitor the developed models and make sure that the models are maintained consistently according to the policies and regulations as new data comes in. This will hold people accountable for their role in model maintenance. Note that the accountability pillar of responsible AI holds people accountable for any model that they create.

## 3. Governance

Data governance is the overall management of the availability, usability, integrity, and security of data used in an enterprise. In broad terms, data governance will take care of the following:

**Data integrity:** According to the DAMA Dictionary of Data Management, data integrity is defined as "data that complies with all rules regarding definitions, relationships, lineages, and heritage" [5]. In other words, data integrity is the practice of preserving the accuracy, consistency, and usability (quality) of the data across different systems.

**Data lineage:** Data lineage is "a description of the pathway from the data source to their current location and the alterations made to the data along that pathway" [5]. Also, tracking the models that use data is part of data lineage when it is considered in the context of AI and machine learning.

**Data security:** Data security is the practice of preserving "…the safety of data from unauthorized and inappropriate access or change. The measures taken to prevent unauthorized access, use, modification, or destruction of data" [5].

**Data loss prevention:** Data loss prevention is "the identification and monitoring of sensitive data to ensure that it's only accessed by authorized users and that there are safeguards against data leaks" [6]. In other words, data loss prevention guarantees no misuse of data.

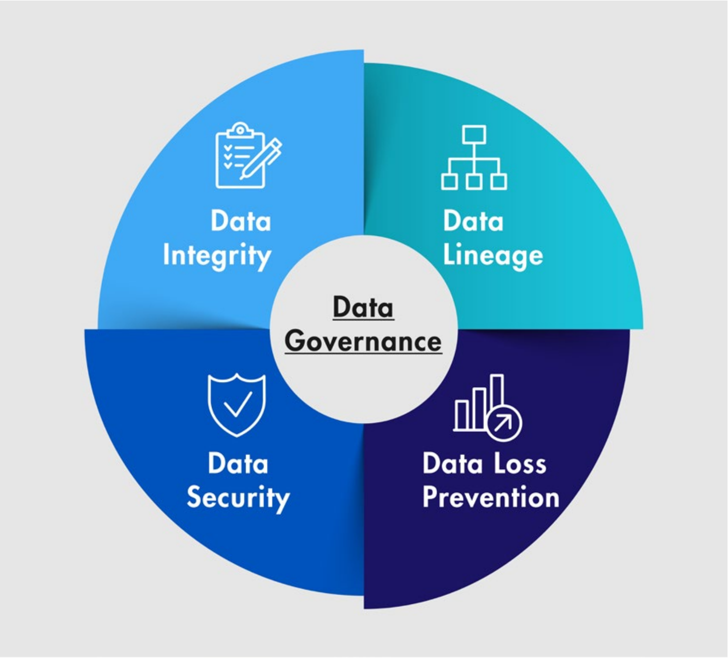

*Figure 2. The components of data governance*

However, data governance for AI is not limited to governing the data handling; one should not think that once data of adequate quality are fed to an AI method, nothing can go wrong with the data. In Figure 3, some standard basics steps for an AI or machine learning algorithm are shown. Data governance must be present in every step.

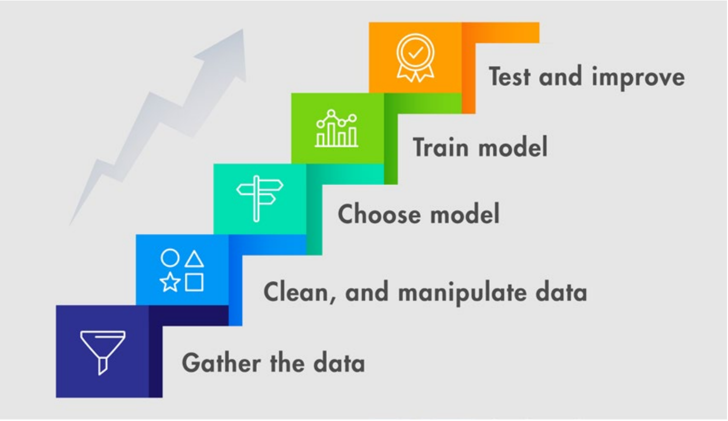

*Figure 3. Steps for applying data to an AI learning algorithm*

Therefore, even though data quality is the foundation of many of these steps and guaranteed by data governance practices, responsible AI requires data governance on the following other layers:

- **Semantic layer.** Data governance on this layer covers how data are labeled and cleaned, what metadata are stored, and what data glossary is used.
- **Modeling and analytical layer.** Data governance for this layer covers subjects such as what modeling techniques are used and how different data are combined.
- **Collection layer.** Data governance needs to guarantee that the data collection follows regulations and makes sure that the data owners have full control over their data and know how their data will be used. Also, the data collection needs to avoid biases.

Therefore, a central data governance framework should be able to know what teams are using data, how, and in which models. Moreover, data collection processes become important in data governance for AI; whoever uses the data should be able to trust that the collected data are reliable, collected under regulation, and free of biases. In general, bias is defined as "distortion of data to support a particular view," and it can exist in data analysis, collection, sampling, and use [5]. It is important to make sure that the AI-related data governance blocks are avoiding biases. For instance, bias can appear when data are collected, used, or labeled to prove a predefined result. Moreover, it is common practice in machine learning to downsample data for some application, a practice that can introduce biases in developed models. Bias has been a term mentioned regularly in association with ethical AI, and it even refers to the presence of bias in data without someone deliberately introducing it [7, 8].

Unlike the specific AI-related concerns of data governance, the main building blocks of data governance are well studied for many industries [9]. The main building blocks of a successful data governance program across most of the industries, shown in Figure 4, are:

## 1. Data stewardship

Successful data governance programs are owned by the business, with strong IT support. Data governance is best implemented by interweaving it in the fabric of the organization by implementing the data stewardship framework. It provides the ability to establish and maintain a clear picture of data ownership within the organization.

## 2. Metadata management

Metadata are "data about data," and they generally define the content of a data object. Metadata provide the means for identifying, defining, and classifying data within the subject areas. The effective management of metadata is one of the essential activities of a data steward within the data governance practice.

## 3. Policy enforcement

A data governance policy is a documented set of guidelines for ensuring the proper management of an organization's digital information. Such guidelines can involve policies for business process management, enterprise risk planning, security, data quality, and privacy. A data governance policy formally outlines how business activity monitoring should be carried out to ensure that organizational data are accurate, accessible, consistent, and protected.

## 4. Data quality

Data quality and data governance are usually interchangeable terms. Data quality is a very important aspect of data governance in establishing consistency, correctness, completeness, and timeliness of the data. Lack of data quality will result in GIGO (garbage in, garbage out).

## 5. Security and privacy

Another important aspect of data governance is the implementation of the security and privacy of data by classifying the data into different levels of security and defining the roles and privileges for those levels of data.

## 6. Change control management

As the saying goes, old habits die hard. It is important to sustain the data governance program once all of the functions are implemented. This is best achieved by introducing a formal change control process.

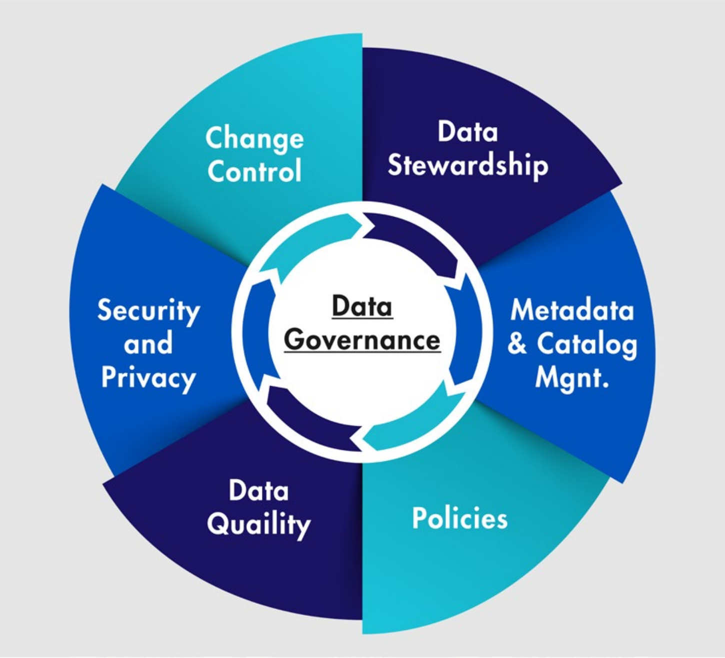

*Figure 4. Main building blocks of a successful data governance program*

## Existing Standards and Policies

The ethics of AI has gained much interest recently, and many companies and governments are working to create definitions, policies, and regulations around the ethics. These policies are stated regardless of the industry of concern and they give a clear roadmap for some industries. However, other industries need to revisit the guidelines in order to best suit their needs.

Many of the efforts described here are in their very first phases and do not pose binding requirements. The following is an overview of some of these activities:

- The European Commission published Ethics Guidelines for Trustworthy AI in April 2019 [10]. The guideline defines trustworthy AI as lawful, ethical, and robust AI. It provides general guidelines to achieve a trustworthy AI, two of which privacy and data governance. This guideline requires that the quality and integrity of data be preserved and that legitimized access to data be guaranteed.
- IEEE's global initiative for ethically aligned design of autonomous and intelligent systems has been working to define standards and guidelines for ethical AI since 2015. These guidelines are required to cover issues such as human rights, accountability, transparency, and awareness of misuse [11]. These standards have addressed standards on system design, transparency of autonomous systems, and data privacy, among others [12].
- The International Standards Organization (ISO) created a new technical subcommittee for AI in 2017 [13]. This subcommittee, SC 42, is working on creating standards as well as addressing issues such as trustworthy AI, safety, and AI applications.

## Responsible AI for the Power Industry

Applications of AI in the power industry continue to grow, but the industry lacks the required standards and policies for responsible AI. Among the three pillars of responsible AI, governance has attracted more attention in this industry. This section summarizes some best practices and considerations associated with this pillar. These practices and considerations are around the main building blocks of data governance (see Figure 4) as well as data collection and labeling.

## Prespectives

### Common Practices for Data Governance in the Power Industry

**Data collection:** Given the variability and massive amount of data sources in the power industry, there is no universal policy or standard for data collection. However, data collection practices in the power industry are mainly dictated by the rules and policies of the owner/provider of the data, such as utilities, and sticking to the most restrictive policies is a common practice. These practices include using similar equipment and configurations for data collection. In some cases, the data collection procedure must comply with policies enforced by entities other than the industry, such as when permits are necessary for flying drones to capture images of transmission assets. Furthermore, states and public service commissions' view data ownership differently. Some consider the utilities to be the owner, whereas others view it to be the customer, which makes it difficult to inform the "owners" how their information is being used.

**Data labeling:** The power industry is in very early stages when it comes to data labeling. The industry usually follows the data provider/owner's preference in naming the data files and labeling the data. However, even for reporting timestamps in time-series data, there is no universal practice that the utilities follow. On the other hand, most of the labeling in the industry is done in a non-automated manner by subject matter experts, which results in different ways and preferences in labeling the data. Moreover, some of the time-series data generated in the industry have very high resolution (thousands of samples per second), which makes it very challenging to label them.

**Data integrity:** Every other organization usually has its own built-in data integrity. In the power industry, data integrity is usually handled based on the data users' policies and may differ based on the project. User authentication, access permission, and code versioning are some of the methods that help with data integrity. However, tracking how data are being combined for different purposes is a big challenge.

**Data lineage:** A standardized information gathering (SIG) questionnaire is a common practice that is followed by many utility companies to ensure data lineage. Also, many utility companies have groups that maintain a database (such as Advanced Metering Infrastructure (AMI) data) and grant user permission so that they can track who has access to the data. However, these practices do not fully cover the types of models that use the data.

**Data security:** Given all the new methods that use data and new practices for data sharing, the utilities are becoming more and more aware of the concerns regarding data security. The North American Electric Reliability Corporation (NERC) Critical Infrastructure protection (CIP) [14] and Federal Energy Regulatory Commission (FERC) Critical Energy/Electric Infrastructure Information (CEII) [15] standards are commonly practiced in the industry. Also, the following are some of the methods that the industry uses for ensuring data security:

- Removal of personally identifiable information from the data set.
- At Rest Obfuscation – Changes to locational, temporal, and sequential elements of the data to make it tedious and difficult to make sense of the information.
- At Rest Encryption – Decryption keys should be stored separately and accommodated only via human authorization and action.
- At Rest (never connected to a private or public network) – Valuable data sets should be put online only when there is a reason for a user to access the information. Availability should be intentionally inconvenient.
- Three-factor authentications for all data access; valuable data sets should be protected with a minimum of two authentication approaches and on human curator authorization.
- Online access monitoring for unusual activity – Any important data set should be actively monitored and subject to periodic reporting.

**Data loss prevention:** Frequent backups of a central data server and having a redundant backup are standard methods that can prevent data loss. Also, all data users should clearly understand their responsibility for preventing the data from being lost.

The above-mentioned practices need to be improved to fully use the power of data governance in the power industry. The following section summarizes some of the specific needs of the power industry for data governance.

## The Considerations for Data Governance in the Power Industry

One of the main concerns for the power industry in dealing with data is to stay up to date on rapidly changing data policies for countries and U.S. states. There are about 25 states with laws in place or being worked on concerning data security/privacy/ownership. This necessitates the existence of a central data governance platform that tracks these changes and adopts the guidelines according to the changes. Below are some considerations specific to different aspects of data governance for the power industry.

**Data collection:** Preserving the privacy of the data owners and making the data untraceable in the data collection process are significant concerns and challenges for the industry. The problem is how to balance these factors with a data collection that gives insight and includes enough information for the use cases that are of interest. One practice that can help with this is explicit statements detailing what usage is acceptable for the data (for example, the data can be used for research, but they cannot be sold as part of a product) and required citations. Moreover, standards such as FERC CEII and NERC CIP need to be followed more rigorously by the stakeholder entities. On the engineering side of data collection, more effort needs to be made to make sure that the labels, units, level of confidences, location, source of data, and so on are all captured and documented correctly and clearly. Also, data variation needs to be considered carefully to train models with better performance.

**Data labeling:** One of the challenges of the application of AI in the power industry is the lack of data, in particular, labeled data. Our industry lacks tools and standard labeling processes and terms. Therefore, labeling data should be done manually, which is practically impossible for some data types and makes the labeled data more prone to human error, and that can affect the quality of the developed models. Moreover, data labeling needs to be done with careful consideration of sensitive customer attributes—for instance, how to consider the demographic data in labeling and avoid revealing customer information at the same time. In many industries, avoiding biases in labeling data is deemed to be very important. However, because many types of data in the power industry are associated with physically measured quantities, the bias in labeling data does not seem to be a problem yet.

**Data integrity:** Models and insights gained are only as good as the underlying data. Therefore, each data set and use case will likely have unique requirements for data integrity. Some use cases may allow lower-quality data (for instance, data with some information being masked), whereas other use cases may require higher-quality data. Even though it is always preferable to have better data, other considerations, such as privacy and data access, need to be considered as a tradeoff. However, establishing a consensus around data quality metrics would be beneficial, so that data consumers can expect a minimum level of quality in the data that they analyze. Cloud-based solutions can potentially help with preserving data integrity. However, utilities are very concerned about integrating third-party cloud-based services because of the lack of best practices and regulations. Also, the industry can borrow best practices from other industries, such as having access and version control across companies, vendors, and so on. Another method is to have some checksum (such as SHA-256) of all the data files [16], thereby providing a way for users to confirm that the files are correct. Providing checksums is standard practice in the machine learning community and provides a relatively straightforward way to verify data integrity. This practice might not be possible with all data types, but it should be possible for most.

There are some use cases that the industry needs to combine data sets to get actionable insights from data. For instance, operational data only gives some insights; if you do not have access to the control system, there is limited use. However, because of the lack of understanding as well as required policies, combining data sets is avoided in general in the industry to make sure that privacy and security are preserved across the grid.

**Data lineage:** Preserving customer privacy is one of the factors that makes data lineage extremely important, for instance, when other customer data are provided along with the AMI data. In general, there is a balance needed between 1) requiring too much metadata, constrained file formats, and so on, thereby limiting the number of datasets available, and 2) not requiring any metadata, structure, validation, and so on, thus making many data sets available but limiting the use. The ideal case is that every data set is presented in a way that makes it clear to users what is contained in the data, the primary use case for the data (regression versus classification, for example), and to what extent the data have been validated (and therefore whether a user can "dive right in" to use the data). Moreover, to the extent possible, there should be a standardized categorization of the data to enable efficient searching and discovery by users. In addition, all data sets should be provided with a brief description of the data; references; or a README with more details on the data itself (such as data formatting, variable naming, and units, measured versus calculated data) and references to the source of the data, both the physical origin of the data (location, company, and so on) and citations of critical reports, articles, and so on that used the data (if available). This will help with tracking the source as well as understanding the models that might use the data sets.

**Data security:** Different stakeholders in the industry are becoming more sensitive about data security. Therefore, data security needs to be made as easy as possible so that any entity who has access to data and uses them can follow the security guidelines. On the other hand, more consideration needs to be given to handling raw data. For instance, raw data should never be accessible unless the authorized user is on-site and the data are not available by network connection. Further, security clearance for trained staff may be warranted. Similar to many other industries, such as healthcare, real-time logging of users' interactions with the data sets and automated alerts for violations can help with data security.

**Data loss prevention:** The power industry can follow the best practices of many other sectors for preventing data losses. For instance, all data should be mirrored in at least one different facility/location and stored in file systems designed specifically for high-integrity data storage (such as Z File System (ZFS)). Additionally, data should be exposed to users in a read-only interface and regularly validated for integrity (for example, by checking the checksums of each file through an automated process run daily or weekly).

## Conclusions

AI has a great potential to use a large volume of data streams to help the industry transition to smarter power generation and utilization, save money, improve safety, and protect the environment. However, integration of AI in the power industry requires that different stakeholders—from engineers to analysts to maintenance workers and IT departments—think deeply about responsible adoption of AI in their area and how they can avoid the misuse of data.

An initial step toward adopting responsible AI is using a trustworthy platform designed to take care of the data governance considerations, such as security and privacy, which are among the primary concern areas for many utilities when it comes to AI adoption or even any data-based solution.

Moreover, the industry needs to understand its specific needs so that any data governance solution is tailored to these needs. These particular needs may cover areas such as data labeling and data privacy, which depend heavily on the industry.

The Electric Power Research Institute is interested in collaborating with industry stakeholders and utilities to understand and further a safe and responsible application of AI in the industry.

# References

[1] An Introduction to AI, Its Use Cases and Requirements for the Electric Power Industry. Electric Power Research Institute, Palo Alto, CA: 2019. [3002017143]

[2] https://www.turing.ac.uk/sites/default/files/2019-06/understanding_artificial_intelligence_ethics_and_safety.pdf

[3] https://www.darpa.mil/program/explainable-artificial-intelligence

[4] https://www.brighttalk.com/webcast/17108/354765?utm_campaign=communication_reminder_starting_now_registrants&utm_medium=email&utm_source=brighttalk-transact&utm_content=button

[5] DAMA Dictionary https://dl.acm.org/citation.cfm?id=2018821

[6] https://www.techopedia.com/definition/25115/data-loss-prevention-dlp

[7] https://www.inc.com/guadalupe-gonzalez/amazon-artificial-intelligence-ai-hiring-tool-hr.html

[8] https://towardsdatascience.com/gender-bias-word-embeddings-76d9806a0e17

[9] https://www.mcpressonline.com/analytics-cognitive/business-intelligence/5-building-blocks-of-a-successful-data-governance-program

[10] https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

[11] https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf

[12] https://ethicsinaction.ieee.org

[13] https://www.iso.org/committee/6794475.html

[14] https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx

[15] https://www.ferc.gov/legal/ceii-foia/ceii.asp

[16] https://www.howtogeek.com/363735/what-is-a-checksum-and-why-should-you-care/

3002017792 January 2020

**Electric Power Research Institute**

3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA •
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com