

# Making Neural Networks FAIR

Anna Nguyen<sup>\*1</sup>, Tobias Weller<sup>\*1</sup>, Michael Färber<sup>1</sup>, and York Sure-Vetter<sup>2</sup>

<sup>1</sup> Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

<sup>2</sup> National Research Data Infrastructure (NFDI), Karlsruhe, Germany

**Abstract.** Research on neural networks has gained significant momentum over the past few years. Because training is a resource-intensive process and training data cannot always be made available to everyone, there has been a trend to reuse pre-trained neural networks. As such, neural networks themselves have become research data. In this paper, we first present the neural network ontology *FAIRnets Ontology*, an ontology to make existing neural network models *findable*, *accessible*, *interoperable*, and *reusable* according to the FAIR principles. Our ontology allows us to model neural networks on a meta-level in a structured way, including the representation of all network layers and their characteristics. Secondly, we have modeled over 18,400 neural networks from GitHub based on this ontology, which we provide to the public as a knowledge graph called *FAIRnets*, ready to be used for recommending suitable neural networks to data scientists.

## 1 Introduction

Researchers of various sciences and data analysts reuse but also re-train neural network models according to their needs.<sup>3</sup> Providing pre-trained neural network models online has the following advantages. First, as a provider, you can benefit from users improving your neural network and circulating your research. Second, as a user of an already trained neural network, you can overcome the cold start problem as well as save on training time and costs. Furthermore, providing trained neural network models gets increasingly important in the light of the research community efforts to make research results more transparent and explainable (see FAIR principles [21]). As a result, more and more trained models are provided online at source code repositories such as GitHub. The models provided serve not only to reproduce the results but also to interpret them (e.g., by comparing similar neural network models). Lastly, providing and using pre-trained models gets increasingly important via transfer learning in other domains.

To ensure the high-quality reuse of data sets and infrastructure, the *FAIR Guiding Principles for scientific data management and stewardship* [21] have been proposed. These guidelines are designed to make digital assets **F**indable, **A**ccessible, **I**nteroperable, and **R**e-usable. They have been widely accepted by

---

<sup>\*</sup> These authors contributed equally to the work.

<sup>3</sup> <https://trends.google.com/trends/explore?date=all&q=transfer%20learning>

several scientific communities nowadays (e.g., [22]). Making digital assets FAIR is essential to deal with a data-driven world and thus keeping pace with an increasing volume, complexity, and creation speed of data. So far, the FAIR principles have been mainly applied when providing data sets and code [22,3], but not machine learning models, such as neural network models. In this paper, we bring the FAIR principles to neural networks by (1) proposing a novel schema (i.e., ontology) which enables semantic annotations to enhance the information basis (e.g., for search and reasoning purposes) and (2) representing a wide range of existing neural network models with this schema in a FAIR way. As we outline in Sec. 3.1, extracting metadata from neural networks automatically is a nontrivial task due to heterogeneous code styles, dynamic coding, and varying versioning. The key idea is that the information contained in these networks should be provided according to the FAIR principles. This comprises several steps which not only consist of having identifiers but providing (meta)data in a machine-readable way in order to enable researchers and practitioners (e.g., data scientists) easy access to the data. We facilitate this by using semantic web technologies such as OWL and RDF/RDFS.

Overall, we provide the following contributions:

1. We provide an *ontology*, called FAIRNETS ONTOLOGY, for representing neural networks. It is made available using a persistent URI by w3id and registered at the platform Linked Open Vocabularies (LOV).

**Ontology URI:** <https://w3id.org/nno/ontology>

**LOV:** <https://lov.linkeddata.es/dataset/lov/vocabs/nno>

2. We provide a *knowledge graph*, called FAIRNETS, representing over 18,400 publicly available neural networks, following the FAIR principles. FAIRNETS is available using a persistent URI by w3id and is uploaded to Zenodo.

**Knowledge Graph URI:** <https://w3id.org/nno/data>

**Zenodo:** <https://doi.org/10.5281/zenodo.3885249>

Our contribution is beneficial in several application areas (see Sec. 5). For instance, we already provide an online search system called FAIRNETS SEARCH [13] by which users can explore and analyze neural network models.

The paper is structured as follows. Sec. 2 describes the structure of FAIRNETS ONTOLOGY and Sec. 3 describes the knowledge graph FAIRNETS. Sec. 4 explains the reason why the neural networks in FAIRNETS follow the FAIR principles. Sec. 5 describes the impact of FAIRNETS. Sec. 6 gives an overview of related work. Lastly, the contributions are summarized.

## 2 FAIRnets Ontology

### 2.1 Creation Process

The FAIRNETS ONTOLOGY is dedicated to model metadata for neural network models on a schema level. We developed the ontology by using Protégé [12]. To

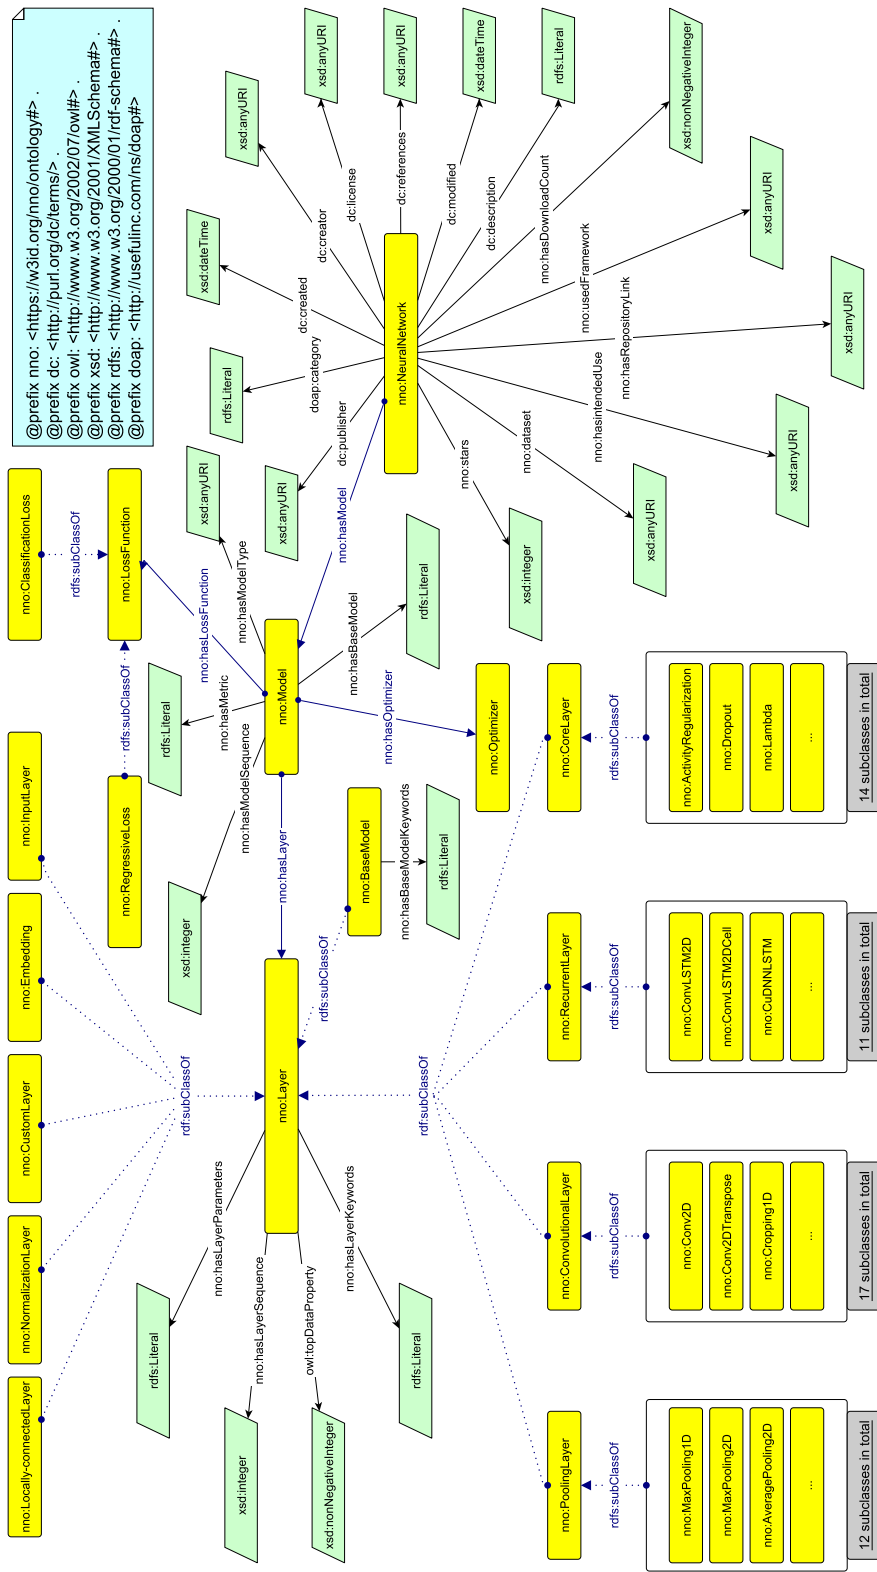


Fig. 1. Visualization of FAIRNETS ONTOLOGY.

the best of our knowledge, there is no existing vocabulary for the specific description of neural networks. That is why several senior researchers use best practices [6] to construct the ontology. We identify researchers, especially beginners, as potential users. The use cases we envision can be found in Sec. 5.

In addition to the consideration of the Predictive Model Markup Language (PMML) in the development of the ontology (especially in describing the architecture), findings from further work were also considered. In particular, model cards [11] were taken into account to validate relevant concepts. Model cards encourage transparent model reports concerning machine learning models and are used for outlining the intended use of a model. These cards define minimal information needed to sufficiently describe a machine learning model (in our case, a neural network) that is relevant to the intended application domains. As suggested from model cards, we included model details such as person developing model, model date, model type, and licenses.

*Characteristics.* The structure of the FAIRNETS ONTOLOGY can be seen in Fig. 1. Overall, the ontology consists of a total of 516 axioms and uses a total of 77 classes where 70 are sub-classes. It also consists of four object properties, 23 data properties, and 29 individuals.

The ontology enables representing three different aspects of information. (1) Neural network-related general metadata and (2) neural network-dependent features can be modeled, such as the type of layer, loss function, and optimizer. (3) Layer-specific metadata is used to enhance the information basis of the specific layers, e.g., its keywords and parameters. In the following, we will describe these three components of the FAIRNETS ONTOLOGY correspondingly.<sup>4</sup>

**General information** describe general components of the neural network, as well as the intended use. For instance, the owner/developer of the (trained) neural network is modeled by using the property `dc:creator`. This attribute makes it possible to search for repositories by the author in the domain of neural networks. Following the Linked Data Principles, the author is represented via a URI. In this way, the authors are uniquely identified. Therefore, it is possible to link it to the Microsoft Academic Knowledge Graph [5] which models scholarly data such as scientific publications in which some of the represented neural network models are proposed and evaluated. Moreover, a name (`rdfs:label`) and a description (`dc:description`) of the trained neural network are stored. The data property `nno:dataset` of type URI allows us to specify the data set that was used to train the neural network. This information already gives a more detailed insight into the neural network as well as the intended use of it.

Furthermore, the timestamp of creation date (`dc:created`) or last modification (`dc:modified`) allows assessing the currency of the neural network. `dc:license` indicates the rights to modify and redistribute that network. Besides, the property `nno:hasRepositoryLink` allows linking to the repository in which the neural network is located. Likewise, references to published papers can be included using `dc:references`.

<sup>4</sup> We will use `nno` as the prefix for the namespace <https://w3id.org/nno/ontology#>

**Model-specific information** covers model-specific components of the neural network, such as optimization function denoted by `nno:hasOptimizer`. The ontology covers modeling various loss functions, such as binary cross-entropy and mean squared error, via the property `nno:hasLossFunction`. Loss functions are subdivided into classification and regression loss functions in the ontology to further indicate the intended use of the neural network. The information about existing layers of the neural network can be linked via the property `nno:hasLayer`. The loss functions and layer types available in Keras, an open-source deep learning framework to model neural networks, served as a basis to model available loss functions and layers.

**Layer-specific metadata** outline additional information about the individual layer. The layers of neural networks are subdivided into subclasses such as core, recurrent, and convolutional layer. These classes are further subdivided into more specific layer classes. This specification derived from Keras enables to categorize the neural networks. For example, a neural network with a layer from class convolutional layer can be assigned to the type convolutional neural network. Furthermore, the hyperparameters (e.g., kernel size, stride, and padding) are denoted by `nno:hasLayerKeywords` and saved as a dictionary. Additional values in the layer are denoted by `nno:hasLayerParameter`.

Most of the categories, properties, and instances are annotated with a label (`rdfs:label`), a description (`rdfs:comment`), and, if given, a link (`rdfs:seeAlso`) which make it easy for ontology users to identify the intended use of categories, properties, and instances, therefore supporting the reusability.

## 2.2 Provisioning

The World Wide Web Consortium (W3C) Permanent Identifier Community Group service is used to provide secure and permanent URL forwarding to the ontology. The FAIRNETS ONTOLOGY in syntax turtle is accessible under <https://w3id.org/nno/ontology>. Moreover, the ontology has been registered at LOV.<sup>5</sup> The ontology is licensed under Creative Commons BY 4.0<sup>6</sup> which allows its wide usage. Furthermore, the ontology follows the 5-Star Linked Data Principles<sup>7</sup> and can, therefore, be easily reused. A VoID file is provided under <https://w3id.org/nno/fairnetsvoid> including provisioning information.

## 3 FAIRnets Knowledge Graph

Apart from the ontology, we provide the FAIRNETS knowledge graph, which is based on the FAIRNETS ONTOLOGY. The knowledge graph allows us to store knowledge (in our case, detailed metadata for neural network models) intuitively as a graph. Existing and widely used W3C standards and recommendations, such as RDF and SPARQL, can be used to query the knowledge graph and to

<sup>5</sup> <https://lov.linkeddata.es/dataset/lov/vocabs/nno>, last acc. 2020-10-15

<sup>6</sup> <https://creativecommons.org/licenses/by/4.0/>, last acc. 2020-10-15

<sup>7</sup> <https://5stardata.info/en/>, last acc. 2020-10-15

integrate relatively easily into existing frameworks and systems. For instance, FAIRNETS is already integrated into KBox [10] which is a data management framework, allowing users to share resources among different applications.

### 3.1 Creation Process

The previous online available neural network repositories such as Keras,<sup>8</sup> Caffe Model Zoo,<sup>9</sup> and Wolfram Alpha<sup>10</sup> are rather small (under one hundred neural networks) and not sufficient to present trends in the development and usage of neural networks. General-purpose online code-sharing services, such as GitHub<sup>11</sup> and Bitbucket,<sup>12</sup> in contrast, contain many repositories of different nature. We, thus, decided to use GitHub since it is the largest host of repositories. Details about the nontrivial extraction process are given in the following.

*Data Source.* We extract and represent metadata of publicly available, trained neural network models in RDF\* (i.e. RDF and RDFS) based on the FAIRNETS ONTOLOGY. Information from SemanGit [9] and GHTorrent [7] can be used to identify GitHub repositories. SemanGit and GHTorrent provide a collection of data extracted from GitHub. In total there are more than 119 million repositories available in the GHTorrent data collection. However, SemanGit and GHTorrent have a different focus and do not provide all the information which we wanted to provide in the FAIRNETS knowledge graph. For instance, information about the architectures of neural networks within the repositories, the creation date of the repositories, as well as the watcher count is not included. We, therefore, directly accessed the GitHub Repository API and queried available neural network repositories. We used the search term ‘neural network’ and filtered for repositories that use Python as a programming language. We accessed these repositories<sup>13</sup> and extracted the neural network metadata.

*Extraction Process.* The difficulty lies in the extraction of the architecture information from the code. We narrowed our extraction down on neural networks implemented in Python. Still, it is difficult to identify the Python file which models a neural network. Therefore, we started with h5 files which are an open-source technology for storing trained machine learning models. Neural networks that have been trained with Keras, for example, can be stored in this format. The h5 file contains information about the neural network, e.g., the sequence of layers, used activation functions, optimization function, and loss function. Accessing the information in the h5 file makes it easier to identify and extract the architecture of the neural network. However, not every repository contains

<sup>8</sup> <https://keras.io>, last acc. 2020-10-15

<sup>9</sup> <https://github.com/BVLC/caffe/wiki/Model-Zoo>, last acc. 2020-10-15

<sup>10</sup> <https://resources.wolframcloud.com/NeuralNetRepository>, last acc. 2020-10-15

<sup>11</sup> <https://www.github.com>, last acc. 2020-10-15

<sup>12</sup> <https://www.bitbucket.org>, last acc. 2020-10-15

<sup>13</sup> Exemplary GitHub API Request: <https://api.github.com/repos/dmnelson/sentiment-analysis-imdb>, last acc. 2020-10-15.

**Table 1.** Mapping of GitHub REST API values to the general components in FAIRNETS.

GitHub API	FAIRnets Ontology
created_at	dc:created
description, readme	dc:description
html_url	nno:hasRepositoryLink
license	dc:license
owner['html_url']	dc:creator
updated_at	dc:modified
watchers_count	nno:stars
name	rdfs:label
topics['names']	doap:category

trained neural networks in h5 files. The reason is that trained neural networks often take up a lot of storage space. Thus, our contribution is the information extraction from the code directly which will be described below.

**General Information:** The mapping of the values from the Github API with the corresponding general component properties in FAIRNETS can be seen in Tab. 1. We use the *full\_name* of the GitHub REST API as a unique identifier (e.g., ‘dmnelson/sentiment-analysis-imdb’ in note 13). The *full\_name* consists of the GitHub username combined with the name of the repository. The owner of the repository is also the owner of the neural network. Moreover, we store the link (**nno:hasRepositoryLink**), the time of creation (**dc:created**), and the last modification (**dc:modified**) of the repository. As a description of the neural network (**dc:description**), we extracted automatically the description and readme file of the GitHub repository. This gives a summary of the possible use of the neural network. Furthermore, license information about the neural network is extracted and modeled in the knowledge graph, if available. This information is crucial regarding the reusability of neural networks. Given this information, it is possible to filter neural networks by license – which is often an important constraint in industrial settings. To enrich the knowledge graph FAIRNETS with information according to the usage of a neural network, we extract the topics<sup>14</sup> of each repository from the GitHub repositories and store them as **doap:category**.

Additionally, we extract arXiv HTTP links within the readme file and map them to **dc:references**. If BibTex file codes can be found in the readme file, we extract the URL information from the BibTex entry and link it by using the property **dc:references**. The property **dc:references** is only intended for scientific contributions. By linking it with URLs from BibTex entries and arXiv links, we ensure this condition. Other links in the readme file are linked to the neural network using **rdfs:seeAlso**.

**Model & Technical Information:** The main feature of FAIRNETS is the modeling of neural networks. We can model the structure and technical compo-

<sup>14</sup> <https://developer.github.com/v3/repos/#list-all-topics-for-a-repository>, last acc. 2020-10-15.

nents of neural networks by employing the FAIRNETS ONTOLOGY. To extract the neural network information from the repositories we consider all Python files in the repositories. Each repository can contain several models of a neural network. In general, it is difficult to extract the architecture information automatically without executing the source code. By executing the code, you can save the neural network model, for example in h5, and retrieve the information easier. We seek a more elegant way by saving execution costs and use language processing to extract the information. Due to that, we focus on Python files with static variables. Despite this restriction, there are still challenges because of various programming styles such as inconsistent naming of variables, complex loop constructions, different structures of code, and other logic statements. Another challenge is changing parameter naming due to different framework versions which are usually not stated. To solve these tasks, a general method is generated using Python Abstract Syntax Trees (AST) module.<sup>15</sup> The AST module helps Python applications to process trees of the Python abstract syntax grammar. We focused on Keras applications of neural networks to extract the architecture because it is the most used deep learning framework among the top-winners on Kaggle.<sup>8</sup> The information on the architecture of the neural network is then modeled by using the schema and properties provided by the FAIRNETS ONTOLOGY. Also, the individual layers and their hyperparameters are stored in our knowledge graph. Likewise, the used optimization function and loss function are stored, among other things, allowing us to infer whether the neural network is used for classification or regression. Our code can be found on GitHub.<sup>16</sup>

*Evaluation.* To evaluate the accuracy of our information extraction, we manually went through 50 examples where we judged the extraction of the GitHub Repository API in Tab. 1. The evaluation was in all cases correct. In the case of the neural network architecture, we used the h5 files, if available, in the repositories. We were able to evaluate over 1,343 h5 files with architecture information (i.e., layer information) which overlap with the architecture extracted from the code with 54% accuracy. Due to later modifications in the code it is possible that the overlap with the h5 file does not apply anymore (e.g., if a layer is commented out).

### 3.2 Provisioning

Just like the FAIRNETS ONTOLOGY, the knowledge graph FAIRNETS is also based on the 5-Star Linked Data Principles. The knowledge graph is accessible under a persistent URI from w3id and additionally provided on Zenodo. In combining FAIR principles and Linked Data Principles using URIs to identify things, providing information using RDF\*, and linking to other URIs, it is possible to easily reference and use FAIRNETS (see Sec. 5). Machine-readable metadata allows us to describe and search for neural networks. The knowledge graph

<sup>15</sup> <https://docs.python.org/3/library/ast.html>, last acc. 2020-10-15

<sup>16</sup> <https://github.com/annnguyen/FAIRnets>



**Table 2.** Statistical key figures about FAIRNETS.

Key Figure	Value
repositories	9,516
unique users	8,637
neural networks	18,463
FFNN	8,924 (48%)
CNN	6,667 (36%)
RNN	2,872 (16%)

FAIRNETS, like the ONTOLOGY, is published under the Creative Commons BY 4.0 license.<sup>6</sup> A VoID file describing the knowledge graph in a machine-readable format is provided under <https://w3id.org/nno/fairnetsvoid>.

### 3.3 Statistical Analysis of the FAIRnets Knowledge Graph

Tab. 2 shows some key figures about the created knowledge graph. It consists of 18,463 neural networks, retrieved from 9,516 repositories, and provided by 8,637 unique users. The creation time of the neural networks in our knowledge graph ranges from January 2015 to June 2019. All these networks have a link to the respective repository and owner. Based on the used layers, we can infer the type of neural network. If a network uses a convolutional layer, it is inferred that the network is a convolutional neural network (CNN). Likewise, if a network contains a recurrent layer, it is inferred that the network is a recurrent neural network (RNN). For simplicity, if none of those two layer types are used, the default claim for the network is a feed-forward neural network (FFNN). Of the total 18,463 neural networks, FFNN is most represented in the knowledge graph comprising half of the neural networks. CNNs follows with 36% and RNN with 16% of the total number of neural networks.

## 4 FAIR Principles for Neural Networks

With FAIRNETS, we treat neural networks as research data. As such, to ensure good scientific practice, it should be provided according to the FAIR principles, that is, the data should be findable, accessible, interoperable, and reusable. While the GitHub repositories themselves do not satisfy the FAIR principles (e.g., the metadata is not easily searchable and processable by machines), the modeling of the neural networks in the FAIRNETS knowledge graph is made *FAIR* as we show in the following. Specifically, in this section, we identify the factors that make the neural network representations in FAIRNETS *FAIR*. This was achieved by following the *FAIRification process*.<sup>17</sup> Our FAIRification process is aligned with

<sup>17</sup> <https://www.go-fair.org/fair-principles/fairification-process/>, last acc. 2020-10-15.

**Table 3.** Evaluation of FAIRNETS according to the Generation2 FAIRMetrics (Note: ✓= passed, (✓)= should pass, ✗= not passed).

Principle	FAIRMetric	Name	Result
Findable	<i>Gen2_FM_F1A</i>	Identifier Uniqueness	✓
	<i>Gen2_FM_F1B</i>	Identifier persistence	✓
	<i>Gen2_FM_F2</i>	Machine-readability of metadata	✓
	<i>Gen2_FM_F3</i>	Resource Identifier in Metadata	✓
	<i>Gen2_FM_F4</i>	Indexed in a searchable resource	✗
Accessible	<i>Gen2_FM_A1.1</i>	Access Protocol	✓
	<i>Gen2_FM_A1.2</i>	Access authorization	✓
	<i>Gen2_FM_A2</i>	Metadata Longevity	(✓)
Interoperable	<i>Gen2_FM_I1</i>	Use a Knowledge Representation Language	✓
	<i>Gen2_FM_I2</i>	Use FAIR Vocabularies	✓
	<i>Gen2_FM_I3</i>	Use Qualified References	✓
Reusable	<i>Gen2_FM_R1.1</i>	Accessible Usage License	✓
	<i>Gen2_FM_R1.2</i>	Detailed Provenance	(✓)
	<i>Gen2_FM_R1.3</i>	Meets Community Standards	(✓)

the *FAIRMetrics*<sup>18</sup> outlined in Tab. 3. In the following, we point out how the single FAIR metrics are met by our knowledge graph.

**Findable** describes the property that metadata for digital assets is easy for both humans and machines to find. Our approach ensured that, firstly, by retrieving the metadata available in the repository, secondly, structuring its metadata in the readme file, and thirdly, obtaining the architecture information from the code file according to the FAIRNETS ONTOLOGY. The neural networks we model have unique identifiers (i.e., fulfilling *Gen2\_FM\_F1A*) and a persistent URI (*Gen2\_FM\_F1B*). As a result, the process for a human to find a suitable neural network through resource identifiers in the metadata (*Gen2\_FM\_F3*) is improved. By using RDF as the data model and by providing a schema in OWL as well as a VoID file as a description of the knowledge graph, the metadata is machine-readable (*Gen2\_FM\_F2*). Thus, the knowledge graph can be automatically filtered and used by services. An exemplary service supporting this statement is presented in Sec. 5. FAIRNETS allows for querying information about and within the architecture of the neural networks which was not possible previously. Now, complex queries are feasible (e.g., list all recurrent neural networks published in 2018), which cannot be solved by traditional keyword searches. The metric *Gen2\_FM\_F4*<sup>19</sup> – ‘indexed in a searchable resource’ – was not passed by FAIRNETS although we indexed it on Zenodo. The reason is that the resource on Zenodo is not findable in the search engine Bing which the au-

<sup>18</sup> <https://fairmetrics.org>, last acc. 2020-10-15

<sup>19</sup> [https://github.com/FAIRMetrics/Metrics/blob/master/FM\\_F4](https://github.com/FAIRMetrics/Metrics/blob/master/FM_F4), last acc. 2020-10-

thors of the FAIRMetrics use as ground truth. However, FAIRNETS is indexed by the search engine Google.

**Accessible** describes that users can access (meta)data using a standardized communication protocol. The protocol must be open, free, and universally implemented. FAIRNETS ONTOLOGY AND KNOWLEDGE GRAPH is located on a web server and can be accessed using the HTTPS protocol (*Gen2\_FM\_A1.1*). The neural networks in the repositories can also be accessed using the HTTPS protocol (*Gen2\_FM\_A1.2*). In addition to the open protocol, the accessible property requires that metadata can be retrieved, even if the actual digital assets are no longer available. Due to the separation of the information in FAIRNETS and the actual neural networks on GitHub, this property is fulfilled, since the information in FAIRNETS is preserved even if the neural networks on GitHub are no longer available (*Gen2\_FM\_A2*). The service to evaluate the metric *Gen2\_FM\_A2* – ‘metadata longevity’ – could not be executed because it only tests files that are less than 300kb<sup>20</sup> whereas FAIRNETS has more than 80MB. This test checks for the existence of the ‘persistence policy’ predicate. This predicate is available in FAIRNETS, which should pass the test.

**Interoperable** refers to the capability of being integrated with other data as well as being available to applications for analysis, storage, and further processing. We make use of Linked Data by applying RDF (*Gen2\_FM\_I1*) and SPARQL to represent the information. This makes the data machine-readable, even without the specification of an ad-hoc algorithm or mapping. Additionally, the FAIRNETS ONTOLOGY and the respective KNOWLEDGE GRAPH use well-established and commonly used vocabularies to represent the information. Among others, Dublin Core, Vocabulary of a Friend (VOAF), Creative Commons (CC), and a vocabulary for annotating vocabulary descriptions (VANN) are used for annotations and descriptions (*Gen2\_FM\_I2*). As a further requirement of the FAIR guideline, qualified references to further metadata are required. This requirement is fulfilled by `rdfs:seeAlso` and `dc:references` (*Gen2\_FM\_I3*). `dc:references` statements provide scientific references between the neural networks and scientific contributions. These references to the scientific contributions are provided via globally unique and persistent identifiers, such as DOIs.

**Reusable** aims at achieving well-defined digital assets. This facilitates the replicability and usage in other contexts (i.e., reproducibility), as well as findability. Due to the architecture and metadata extraction, the process of finding and reusing a neural network by an end-user becomes significantly easier and can now be performed systematically. By using best practices in ontology building, the properties and classes of FAIRNETS ONTOLOGY provided are self-explanatory with labels and descriptions (*Gen2\_FM\_R1.3*). The neural networks in FAIRNETS contain structured detailed metadata such as creator and GitHub link (see *Gen2\_FM\_R1.2*) for easy findability and reuse. At the same time, most neural networks in FAIRNETS have an assigned license which is important for reusability (*Gen2\_FM\_R1.1*). For passing *Gen2\_FM\_R1.2*, (meta)data must be associ-

<sup>20</sup> [https://github.com/FAIRMetrics/Metrics/blob/master/FM\\_A2](https://github.com/FAIRMetrics/Metrics/blob/master/FM_A2), last acc. 2020-10-

ated with detailed provenance reusing existing vocabularies such as Dublin Core which we included in our knowledge graph. *Gen2\_FM\_R1.3* tests a certification saying that the resource is compliant with a minimum of metadata. FAIRNETS is described by using LOV standards for publication. Therefore, we assume that these metrics are fulfilled. Overall, the neural networks modeled in FAIRNETS fulfill all requirements of the FAIR principles, see Tab. 3.

## 5 Impact

We see high potential of FAIRNETS ONTOLOGY and the knowledge graph FAIRNETS in the areas of *transparency*, *recommendation*, *reusability*, *education*, and *search*. In the following, we outline these application areas in more detail.

*Transparency.* Neural networks are applied in many different areas such as finance [16], medical health [8], and law [14]. Transparency plays a major role in these areas when it comes to trust the output of a used neural network model. We claim that our contribution which makes neural networks more transparent can increase trust and privacy [17]. Additionally, using semantic annotations can even enhance interpretability by distributional semantics [18].

Another aspect is the transparency of scientific work regarding neural networks. Researchers publishing a model should provide it according to the FAIR principles to strengthen their scientific contribution. Our knowledge graph FAIRNETS can pave the way for this.

*Recommendation.* Neural Architecture Search (NAS) is used to find the best suitable neural network architecture based on existing architectures [4]. However, the search is performed purely based on metrics like accuracy ignoring explainability aspects concerning the best fitting model. Our knowledge graph allows us to have a search for the best suitable neural network models on a meta-level, using modeled use-cases, data sets, and scientific papers. Knowledge graphs have also been used to provide explanations for recommendation to the user [19,23].

Additionally, we can apply explainable reasoning [20] given the ontology and the knowledge graph and infer some rules. Doing this, we might reason which neural network models are reasonable or which components of the architecture stand in conflict with each other.

*Reusability.* Transfer learning is a method in deep learning to reuse pre-trained models on new tasks. Our contribution facilitate the search of pre-trained neural networks and provide metadata needed to choose a specific neural network. We can envision FAIRNETS linked with other knowledge bases to enrich reusability of neural networks by applying Linked Data Principles.<sup>7</sup> For example, training data sets can be linked with Neural Data Server,<sup>21</sup> Wikidata,<sup>22</sup> and Zenodo<sup>23</sup>

<sup>21</sup> <http://aidemos.cs.toronto.edu/nds/>, last acc. 2020-10-15

<sup>22</sup> <https://www.wikidata.org>, last acc. 2020-10-15

<sup>23</sup> <https://zenodo.org>, last acc. 2020-10-15

through schema.org,<sup>24</sup> scientific papers can be linked with the Microsoft Academic Knowledge Graph [5], and metadata can be extended with OpenAIRE.<sup>25</sup>

On the other hand, providing a model and encouraging its reuse can improve it by revealing limitations, errors, or suggestions to other tasks.

*Education.* Our FAIRNETS knowledge graph can be used for educational purposes [2], for instance, to learn best practices regarding designing a neural network model. Another aspect is to learn the usages of different architectures and their approaches (e.g., via linked papers). Our knowledge graph includes training parameters that can help setting up the training process of a neural network (e.g., when facing the cold start problem).

*Search* We provide online the search system FAIRNETS SEARCH<sup>26</sup> [13], which is based on the proposed FAIRNETS ONTOLOGY and KNOWLEDGE GRAPH. Users can search for neural network models through search terms. Additional information can be retrieved by using SPARQL as query language on top of our knowledge graph, which enables faceted and semantic search capabilities. The SPARQL endpoint is also available to the public. The search system shows how a semantic search system can be realized which improved the limited capabilities of keyword searches on GitHub. Furthermore, developers can provide their GitHub repository to run the FAIRification process on their neural networks. Until now, we have over 550 visits to the website FAIRNETS SEARCH with over 4,800 page views, 1,400 searches on our website with an average duration of twelve minutes, and the maximal actions in one visit is 356.

## 6 Related Work

*Information of neural network models.* Mitchell et al. [11] suggest which information about neural networks should be considered as relevant when modeling them. Information such as description, date of the last modification, link to papers, or other resources to further information, as well as the intended purpose of a neural network, are taken into account. Storing such information makes the neural networks more transparent. We follow this suggestion by defining a semantic representation which, to the best of our knowledge, does not exist for neural network models so far.

The knowledge extraction from neural networks can point out relevant features or redundancies [1]. We extract neural network information to build a knowledge graph to better evaluate the causal relationships between different neural network architectures.

<sup>24</sup> <https://schema.org>, last acc. 2020-10-15

<sup>25</sup> <https://www.openaire.eu>, last acc. 2020-10-15

<sup>26</sup> <https://km.aifb.kit.edu/services/fairnets/>, last acc. 2020-10-15

*Representing and provisioning neural network models.* There exist several standards for the exchange of neural network information on the instance level. The Predictive Model Markup Language (PMML)<sup>27</sup> is an XML-based standard for analytic models developed by the Data Mining Group. PMML is currently supported by more than 30 organizations. The Open Neural Network eXchange format (ONNX)<sup>28</sup> is a project by Facebook and Microsoft that converts neural networks into different frameworks. These two formats serve as an exchange format for neural networks on instance level. We are less interested in the exchange of formats, but rather the reusability of the neural networks on a meta-level. Therefore, our FAIRNETS ONTOLOGY lifts its elements to a semantic level, i.e. to RDF/S, following a methodology for reusing ontologies [15] and applying the Linked Data Principles.<sup>7</sup> Thus, we incorporate information on the instance and meta-level in the knowledge graph FAIRNETS.

*Neural network repositories.* Many pre-trained neural networks are available online. The well-known Keras framework offers ten pre-trained neural networks for reuse.<sup>8</sup> The Berkeley Artificial Intelligence Research Lab has a deep learning framework called Caffe Model Zoo<sup>9</sup> which consists of about fifty neural networks. Wolfram Alpha has a repository with neural networks<sup>10</sup> which consists of approximately ninety models. These pre-trained neural networks are represented in different formats making it, for instance, difficult to compare or reuse neural networks. Besides, a larger number of neural networks can be found in code repositories such as GitHub. These neural networks are typically coded in one of the major programming frameworks such as Keras, TensorFlow, or PyTorch. Our approach aims to consider such neural networks and make them available as FAIR data.

## 7 Conclusion and Future Work

This paper was dedicated to make neural networks FAIR. To this end, we first proposed the FAIRNETS ONTOLOGY, an ontology that allows us to model neural networks on a fine-grained level and that is easily extensible. Second, we provided the knowledge graph FAIRNETS. This graph contains rich metadata of 18,463 publicly available neural network models using our proposed ontology as knowledge schema. Both the FAIRNETS ONTOLOGY as well as the FAIRNETS knowledge graph show a high potential impact in fields like recommender systems and education.

For the future, we plan to connect the FAIRNETS ONTOLOGY and KNOWLEDGE GRAPH with scholarly data. Specifically, we will work on linking publications, authors, and venues modeled in knowledge graphs like the Microsoft Academic Knowledge Graph [5] or Wikidata to the FAIRNETS knowledge graph. This will require to apply sophisticated information extraction methods on scientific publications.

<sup>27</sup> <http://dmg.org/pmml/v4-0-1/NeuralNetwork.html>, last acc. 2020-10-15

<sup>28</sup> <https://onnx.ai>, last acc. 2020-10-15

## References

1. Boger, Z., Guterman, H.: Knowledge extraction from artificial neural network models. In: Proc. of SMC'97. pp. 3030–3035 (1997)
2. Chen, P., Lu, Y., Zheng, V.W., et al.: KnowEdu: A System to Construct Knowledge Graph for Education. *IEEE Access* **6**, 31553–31563 (2018)
3. Devarakonda, R., Prakash, G., Guntupally, K., et al.: Big Federal Data Centers Implementing FAIR Data Principles: ARM Data Center Example. In: Proc. of BigData'19. pp. 6033–6036 (2019)
4. Elsken, T., Metzen, J.H., Hutter, F.: Neural Architecture Search: A Survey. *J. Mach. Learn. Res.* **20**, 55:1–55:21 (2019)
5. Färber, M.: The Microsoft Academic Knowledge Graph: A Linked Data Source with 8 Billion Triples of Scholarly Data. In: Proc. of ISWC'19. pp. 113–129 (2019)
6. Gangemi, A., Presutti, V.: Ontology Design Patterns. In: Handbook on Ontologies, pp. 221–243. International Handbooks on Information Systems (2009)
7. Gousios, G.: The GHTorrent dataset and tool suite. In: Proc. of MSR'13. pp. 233–236 (2013)
8. Khan, J., Wei, J.S., Ringnér, M., et al.: Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks. *Nature Medicine* **7**(6), 673–679 (2001)
9. Kubitza, D.O., Böckmann, M., Graux, D.: SemanGit: A Linked Dataset from git. In: Proc. of ISWC'19. pp. 215–228 (2019)
10. Marx, E., Soru, T., Baron, C., Coelho, S.A.: KBox: Distributing Ready-to-Query RDF Knowledge Graphs. In: Proc. of ESWC'17. pp. 54–58 (2017)
11. Mitchell, M., Wu, S., Zaldivar, A., et al.: Model Cards for Model Reporting. In: Proc. FAT'19. pp. 220–229 (2019)
12. Musen, M.A.: The Protégé Project: A Look Back and a Look Forward. *AI Matters* **1**(4), 4–12 (2015)
13. Nguyen, A., Weller, T.: FAIRnets Search - A Prototype Search Service to Find Neural Networks. In: Proc. of SEMANTICS'19. vol. 2451 (2019)
14. Palocsay, S.W., Wang, P., Brookshire, R.G.: Predicting criminal recidivism using neural networks. *Socio-Economic Planning Sciences* **34**(4), 271 – 284 (2000)
15. Pinto, H.S., Martins, J.P.: Reusing Ontologies. AAI Technical Report SS-00-03 pp. 77–84 (2000)
16. Qi, S., Jin, K., Li, B., et al.: The exploration of internet finance by using neural network. *J. Comput. Appl. Math.* **369** (2020)
17. Schwabe, D.: Trust and Privacy in Knowledge Graphs. In: Proc. of WWW'19. pp. 722–728 (2019)
18. Silva, V.D.S., Freitas, A., Handschuh, S.: On the Semantic Interpretability of Artificial Intelligence Models. *CoRR* **abs/1907.04105** (2019)
19. Wang, X., He, X., Cao, Y., et al.: KGAT: Knowledge Graph Attention Network for Recommendation. In: Proc. of KDD'19. pp. 950–958 (2019)
20. Wang, X., Wang, D., Xu, C., et al.: Explainable Reasoning over Knowledge Graphs for Recommendation. In: Proc. of AAAI'19. pp. 5329–5336 (2019)
21. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., et al.: The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* **3** (2016)
22. Wise, J., de Barron, A.G., Splendiani, A., et al.: Implementation and relevance of FAIR data principles in biopharmaceutical R&D. *Drug Discovery Today* **24**(4), 933 – 938 (2019)
23. Xian, Y., Fu, Z., Muthukrishnan, S., et al.: Reinforcement Knowledge Graph Reasoning for Explainable Recommendation. In: Proc. SIGIR'19. pp. 285–294 (2019)