



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/191957/>

Version: Submitted Version

Preprint:

Lu, H., Liu, X., Turner, R. et al. (Submitted: 2021) PyKale: knowledge-aware machine learning from multiple sources in Python. [Preprint - arXiv]

<https://doi.org/10.48550/arXiv.2106.09756>

© 2021 The Author(s). For reuse permissions, please contact the Author(s).

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

PyKale: Knowledge-Aware Machine Learning from Multiple Sources in Python

Haiping Lu

The University of Sheffield
Sheffield, United Kingdom
h.lu@sheffield.ac.uk

Xianyuan Liu

The University of Sheffield
Sheffield, United Kingdom
xianyuan.liu@sheffield.ac.uk

Robert Turner

The University of Sheffield
Sheffield, United Kingdom
r.d.turner@sheffield.ac.uk

Peizhen Bai

The University of Sheffield
Sheffield, United Kingdom
PBai2@sheffield.ac.uk

Raivo E Koot

The University of Sheffield
Sheffield, United Kingdom
rekoot1@sheffield.ac.uk

Shuo Zhou

The University of Sheffield
Sheffield, United Kingdom
SZhou20@sheffield.ac.uk

Mustafa Chasmai

Indian Institute of Technology, Delhi
New Delhi, India
cs1190341@iitd.ac.in

Lawrence Schobs

The University of Sheffield
Sheffield, United Kingdom
laschobs1@sheffield.ac.uk

ABSTRACT

Machine learning is a general-purpose technology holding promises for many interdisciplinary research problems. However, significant barriers exist in crossing disciplinary boundaries when most machine learning tools are developed in different areas separately. We present *Pykale* – a Python library for knowledge-aware machine learning on graphs, images, texts, and videos to enable and accelerate interdisciplinary research. We formulate new green machine learning guidelines based on standard software engineering practices and propose a novel *pipeline*-based application programming interface (API). *PyKale* focuses on leveraging knowledge from multiple sources for accurate and interpretable prediction, thus supporting multimodal learning and transfer learning (particularly domain adaptation) with latest deep learning and dimensionality reduction models. We build *PyKale* on PyTorch and leverage the rich PyTorch ecosystem. Our pipeline-based API design enforces standardization and minimalism, embracing *green machine learning* concepts via reducing repetitions and redundancy, reusing existing resources, and recycling learning models across areas. We demonstrate its interdisciplinary nature via examples in bioinformatics, knowledge graph, image/video recognition, and medical imaging.

KEYWORDS

machine learning, deep learning, domain adaptation, multimodal learning, transfer learning

1 INTRODUCTION

Machine learning is the cornerstone technology for artificial intelligence (AI), driving many advances in our everyday lives and industrial sectors. AI research becomes more and more interdisciplinary as many problems rely on expertise from various domains. We have also witnessed many machine learning models transverse different research areas and disciplines. For example, the success of convolutional neural networks (CNNs) [26] has spread from computer vision to graph analysis via graph convolutional networks (GCN) [23] and medical imaging via U-net [49], and transformers

[59] developed in natural language processing (NLP) have become a hot topic in solving vision tasks [6, 10, 60].

With rapid development and growing interests in machine learning, many researchers hope to solve real-world interdisciplinary problems using machine learning. However, even with the popularity of open-source software and high-level scripting language such as Python, navigating the abundant choices and variety of machine learning software is not trivial. Researchers often run into barriers when adapting a machine learning tool for a new task of their interest. Solving complex real-world problems in practice often involve analyzing multiple sources of data, e.g., multiple modalities, multiple domains, and multiple knowledge bases. Most machine learning software packages are developed with a specific domain of application in mind. While popular generic packages such as PyTorch and TensorFlow support multiple domains and are not tailored for a specific domain, their focus on generic frameworks makes them inadequate to directly support interdisciplinary research where both flexible configurations and high-level integration are important.

In this paper, we propose *PyKale*, an open-source Python library to enable and accelerate interdisciplinary research via knowledge-aware multimodal learning and transfer learning on graphs, images, texts, and videos. It aims to fill the gaps between rich data sources, abundant machine learning libraries, and eager interdisciplinary researchers, with a focus on leveraging knowledge from multiple sources for accurate and interpretable prediction. To the best of our knowledge, this is the first publicly available Python library that considers both multimodal learning and transfer learning under a common framework of learning from multiple sources. It will make latest machine learning tools more accessible and accelerate the development of such tools. The name of the library consists of **Py** for Python, and **Kale** for Knowledge-aware learning.

PyKale proposes a novel *pipeline-based* application programming interface (API) to enforce standardization and minimalism, as shown in Figure 1. It advocates our newly formulated **green machine learning** concepts of reducing repetitions and redundancy (Fig. 2(a)), reusing existing resources (Fig. 2(b)), and recycling learning models across areas (Fig. 2(c)) by building on standard software

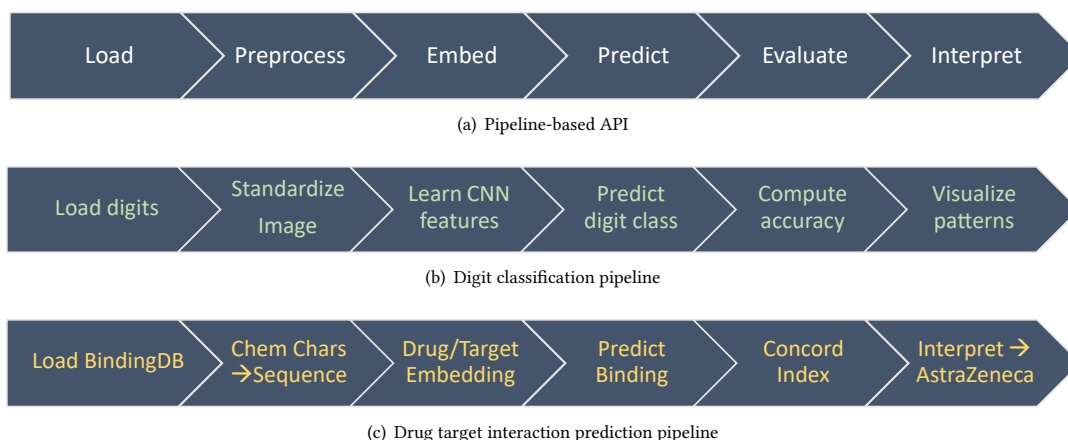


Figure 1: The proposed pipeline-based API in PyKale and two real-world examples.



Figure 2: Green machine learning concepts in PyKale.

engineering practices, extending them, and tailoring the philosophies to machine learning. We include examples in bioinformatics, knowledge graph, image/video recognition, and medical imaging in PyKale. This library was motivated by needs in healthcare applications and thus considers healthcare as a primary domain of usage. PyKale is largely built on PyTorch and leverages many packages in the PyTorch ecosystem,¹ with the aim to become part of it. Our logo in Fig. 2(d) reflects the above characteristics, using an icon of simplified kale leaves.

Specifically, the main contributions are fourfold:

- We introduce the PyKale library for knowledge-aware machine learning. It focuses on leveraging knowledge from multiple sources for accurate and interpretable prediction to enable and accelerate interdisciplinary research.
- We propose a pipeline-based API for a standardized and minimal design to help break interdisciplinary barriers. We advocate green machine learning concepts of reduce, reuse, and recycle in such a design.
- We demonstrate the usage of PyKale on real-world examples from multiple disciplines including bioinformatics, knowledge graph, image/video recognition, and medical imaging.
- We provide many community-engaging features including a detailed documentation, a project board to show the progress and road map, and GitHub discussions open to all users.

The rest of this paper is organized as follows. We first review the state of the art open source software packages that are related

¹<https://pytorch.org/ecosystem/>

to PyKale in Section 2. Then we discuss the design principles and API structure of PyKale in Section 3. Next, we describe the usage of PyKale in Section 4 and show example use cases from different applications in Section 5. Finally, we show the openness of our package and discuss its limitations and future developments in Section 6, with conclusions drawn in Section 7.

The PyKale library is publicly available at <https://github.com/pykale/pykale> with accompanying data (mainly for testing at the moment) at <https://github.com/pykale/data> under an MIT license. PyKale can be installed from the Python Package Index (PyPI) via `pip install pykale`. PyKale documentation is hosted at <https://pykale.readthedocs.io>. The primary targeted users are researchers and practitioners who have experience in Python and PyTorch programming and need to apply or develop machine learning systems taking data from multiple sources for prediction tasks, particularly in interdisciplinary areas such as healthcare. This paper refers to release version **0.1.0rc2**.

2 RELATED WORK

As an open source project, we have learned from numerous libraries in the public domain to build ours. Here, we can only briefly mention several that have been particularly influential or relevant. In particular, we focus on those PyTorch-based libraries that we have frequently studied in our development, regrettably omitting many libraries, such as those based on TensorFlow.

2.1 PyTorch ecosystem

PyTorch is a popular open source machine learning library, particularly for computer vision and NLP applications. The PyTorch ecosystem has a rich collection of tools and libraries for the development of advanced machine learning and AI systems. PyKale aims to fill the gap within the PyTorch ecosystem to support more interdisciplinary research based on multiple data sources. Therefore, we make extensive usage of existing libraries from the PyTorch ecosystem to reduce duplicated implementation.

2.1.1 PyTorch Lightning. PyTorch Lightning is a popular deep learning framework providing a high-level interface for PyTorch. By

removing boilerplate code, it simplifies the development of research code and improves the reproducibility, flexibility, and readability of the resulting models [11]. The goal of PyKale shares some similarity with PyTorch Lightning, but with a different focus on supporting interdisciplinary research. We have lots of inspirations from the design of PyTorch Lightning.

2.1.2 Other PyTorch libraries. PyKale depends on some other libraries from the PyTorch ecosystem including TensorLy for tensor analysis [24], TorchVision for computer vision [37], and PyTorch Geometric for graph analysis [13]. We also learned from MONAI for medical image analysis [36], GPyTorch for Gaussian processes [16], Kornia for computer vision [48], and TorchIO for medical imaging preprocessing [45]. These libraries have focuses different from ours on interdisciplinary research and multiple data sources.

The above libraries are listed as references at the end of our contributing guideline page.²

2.2 Multimodal/transfer learning libraries

To the best of our knowledge, no other libraries in the PyTorch ecosystem have the same pipeline-based API design as PyKale. Several PyTorch-based libraries on multimodal learning or transfer learning are summarized below.

2.2.1 MultiModal Framework (MMF). MMF [52] is the only library in the current PyTorch ecosystem focusing on multimodal learning of vision and language data in applications such as visual question answering, image captioning, visual dialog, hate detection and other vision and language tasks. PyKale differs from MMF not only in the API design, but also in the scientific fields covered and interdisciplinarity. PyKale aims to support interdisciplinary research such as medical imaging and drug discovery, and includes examples in these areas in the current release.

2.2.2 Transfer-Learning-Library. The Transfer-Learning-Library [21] is a library on transfer learning, providing domain adaptation and fine tuning algorithms for computer vision applications. PyKale differs not only in the API design but also in the multiple modalities of data supported, including also graphs and texts, as well as in interdisciplinarity.

2.2.3 Cornac. Cornac [50] is a library for multimodal recommender systems, leveraging auxiliary data (e.g., item descriptive text and image, social network, etc). PyKale differs from Cornac not only in the API design but also in the more diverse machine learning models and applications supported, not limited to recommender systems.

2.2.4 ADA. ADA [56], with a full name “(Yet) Another Domain Adaptation library”, is an excellent package built on PyTorch Lightning for unsupervised and semi-supervised domain adaptation. We refactored ADA and made many changes for adaption to our pipeline-based API. We include a docstring at the top of each module adapted from ADA to indicate the source at ADA. Beyond a new API design, PyKale extends ADA substantially to support video domain adaptation and also supports non-vision data for interdisciplinary research.

There are some other smaller libraries on multimodal or transfer learning that are more narrowly focused than the above, such as Multimodal-Toolkit [17]. PyKale frames multimodal learning and transfer learning under one roof of knowledge-aware machine learning from multiple sources with a unified pipeline-based API, aiming to support interdisciplinary research rather than just popular vision or language tasks.

3 PYKALE DESIGN

3.1 Green machine learning

Green machine learning (and **green AI**) is a scarcely used term referring to energy-efficient computing [5, 15, 27]. Here, we propose a new *green machine learning* perspective for machine learning software development by formulating the **3R** guiding principles below. We build these principles on standard software engineering practices by extending them and tailoring the philosophies to machine learning:

- **Reduce** repetition and redundancy
 - Refactor code to standardize workflow and enforce styles, e.g., we refactored the deep drug-target binding affinity (DeepDTA) [64] into our PyKale pipeline (Fig. 1(c)).
 - Identify and remove duplicated functionalities, e.g., construct data loading API for popular dataset to share among different projects.
- **Reuse** existing resources
 - Reuse the same machine learning pipeline for different data and applications, such as using the same multilinear principal component analysis (MPCA) pipeline for gait [35], brain [53], and heart [55].
 - Reuse existing libraries (e.g., those in the PyTorch ecosystem, such as PyTorch Geometric) for available functionalities rather than implementing them again.
- **Recycle** learning models across areas
 - Identify commonalities between applications, e.g., the similarity between commercial recommender systems (predicting user-item interactions) and drug discovery (predicting drug-target interactions).
 - Recycle models for one application to another, e.g. from recommender systems [2] to drug discovery [61].

Although the above is largely based on standard software engineering practices, this new formulation offers a new perspective to focus on core principles of standardization and minimalism. It has guided us to design a unique pipeline-based API to unify workflow, break barriers between areas and applications, and cross boundaries to fuse existing ideas and nurture new ideas.

3.2 Pipeline-based API

Inspired by the convenience of machine learning pipelines in machine learning library like Spark MLlib [38] and scikit-learn [44], we design PyKale with a pipeline-based API as shown in Fig. 1(a). This design has six key steps and embodies our green machine learning principles above by organizing code along a standardized machine learning pipeline to identify commonalities, reduce redundancy, and minimize cognitive overhead.

²<https://github.com/pykale/pykale/blob/main/.github/CONTRIBUTING.md>

```

1 # Load digits from multiple sources [digits_dann_lightn/main.py]
2 from kale.loaddata.digits_access import DigitDataset
3 from kale.loaddata.multi_domain import MultiDomainDatasets
4
5 source, target, _ = DigitDataset.get_source_target(
6     DigitDataset("MNIST"), DigitDataset("USPS"), data_path)
7 dataset = MultiDomainDatasets(source, target)
8
9 # Preprocess digits [kale/loaddata/digits_access.py]
10 import kale.prepdata.image_transform as image_transform
11
12 self._transform = image_transform.get_transform(transform_kind)
13 def get_train(self):
14     return MNIST(data_path, train=True, transform=self._transform)
15
16 # Embed digit representations [digits_dann_lightn/model.py]
17 from kale.embed.image_cnn import SmallCNNFeature
18
19 # Predict digit class and domain [digits_dann_lightn/model.py]
20 from kale.predict.class_domain_nets import ClassNetSmallImage,
21     DomainNetSmallImage
22
23 # Build domain adaption pipeline [digits_dann_lightn/model.py]
24 import kale.pipeline.domain_adapter as domain_adapter
25
26 model = domain_adapter.create_dann_based(method="DANN",
27     dataset=dataset,
28     feature_extractor= SmallCNNFeature(),
29     task_classifier=ClassNetSmallImage(),
30     critic=DomainNetSmallImage(),
31     **train_params)
32
33 # Utility functions [digits_dann_lightn/main.py]
34 from kale.utils.csv_logger import setup_logger
35 from kale.utils.seed import set_seed

```

Code 1: Code snippets from the source code for the digits domain adaptation example at `pykale/examples/digits_dann_lightn/main.py` to demonstrate the unified pipeline-based API, simplified for inclusion here.

In the following, we explain our unified API by starting with what the input and output are. We provide Python code snippets to help this explanation in Code 1, mainly using the domain adaptation for digit classification example with a pipeline in Fig. 1(b).³ Figure 1(c) shows another pipeline for drug discovery. More code snippets are in Section 5.

3.2.1 Load. The `kale.loaddata` module mainly takes source paths (local or online) as the input and constructs dataloaders for datasets as the output. Its primary function is to load data for input to the machine learning system/pipeline. See line 2–7 of Code 1 for an example of loading digit images from multiple sources and line 8–13 of Code 4 for an example of loading drug and targets data.

3.2.2 Preprocess. The `kale.prepdata` module takes the loaded raw input data as input and preprocesses (transforms) them into a suitable representation for the following machine learning modules.

Preprocessing steps include data normalization, augmentation, and other transformations of data representation not involving machine learning. Its submodules are typically imported in `kale.loaddata`. See line 10–14 of Code 1 for an example of standardizing digit images with predefined transforms.

3.2.3 Embed. The `kale.embed` module takes preprocessed, normalized data representations to *learn* new representations in a new space as the output. It includes dimensionality reduction algorithms (feature extraction and feature selection), such as MPCA [35] and CNNs. They can be viewed as encoders or embedding functions that learn suitable representations from data. This is a machine learning module. See line 17 of Code 1 for an example of (importing) a CNN feature extractor.

3.2.4 Predict. The `kale.predict` module takes the learned (or preprocessed, if skipping `kale.embed`) representations to predict a desired target value as the output. Thus, this module provides prediction functions or decoders that learn a mapping from the input representation to a target prediction. This is also a machine learning module. See line 20–21 of Code 1 for an example of (importing) digit and domain classifiers.

3.2.5 Evaluate. The `kale.evaluate` module evaluates the prediction performance using some metrics. We reuse metrics from other libraries (e.g., `sklearn.metrics` in line 4 of Code 5) and only implement metrics not commonly available, such as the Concordance Index (CI) [1] for measuring the proportion of concordant pairs. See line 19 of Code 4 for its example usage.

3.2.6 Interpret. The `kale.interpret` module aims to provide functions for interpretation of the learned features, the model, or the prediction results/outputs, e.g., via further analysis or visualization, and we only implement functions not commonly available. This module has implemented functions for selecting and visualizing weights from a trained model. See line 16–20 of Code 5 for an example of visualizing weights of a linear model for interpretation.

3.2.7 Pipeline. The `kale.pipeline` module provides mature, off-the-shelf machine learning pipelines for “plug-in usage”. Its submodules typically specify a machine learning workflow by combining several other modules. See line 24–31 of Code 1 for an example of calling a domain adaptation pipeline.

3.2.8 Utilities. The `kale.utils` module provides common utility functions, such as setting random seeds, logging results, or downloading data. See line 34–35 of Code 1 for examples of importing the seed-setting and csv-logging submodules.

3.3 Machine learning models

Machine learning models in PyKale can be categorized into four main (possibly overlapping) groups.

3.3.1 Multimodal learning. To support learning from data of multiple modalities, we need to first support learning from each individual modality. Thanks to the rich PyTorch ecosystem, we can build upon other libraries to have machine learning models supporting graphs, images, texts, and videos, primarily using PyTorch Dataloaders. The only missing major modality is audio but we have ongoing effort to include it in the near future, building upon `torchaudio`.

³https://github.com/pykale/pykale/tree/main/examples/digits_dann_lightn

Learning from heterogeneous data sources and data integration can be viewed as multimodal learning as well. To this end, PyKale has built a DeepDTA [42] pipeline `kale.pipeline.deep_dti` that learns from drug and target data, the chemical representation of which can be transformed into sequence or vector representations. PyKale also implemented our recent Graph information propagation Network (GripNet) [61] `kale.embed.gripnet` for link prediction and data integration on heterogeneous knowledge graphs. PyKale has an example `drug_gripnet` to show the model usage on public bioinformatics knowledge graph with drug and protein's information.

3.3.2 Transfer learning. In transfer learning, PyKale currently focuses on domain adaptation [3, 43]. We largely inherited the excellent, modular architecture from ADA [56], covering many important semi-supervised and unsupervised domain adaptation algorithms, such as domain-adversarial neural networks (DANN) [14], conditional adversarial domain adaptation networks (CDAN) [33], deep adaptation networks (DAN) [32] and joint adaptation networks (JAN) [34], and Wasserstein distance guided representation learning (WDGRL) [51]. These algorithms are applicable to all modalities with appropriate representations. PyKale currently has two pipelines for domain adaptation: `domain_adapter` and `video_domain_adapter` in `kale.pipeline`.

3.3.3 Deep learning. PyKale builds deep neural networks (DNNs) upon the PyTorch API. Current implementations include CNNs [26] / 3D CNNs [7, 57], GCNs [23], and attention-based networks such as transformers [59] and squeeze and excitation networks [18] (see more in Section 5). We use TorchVision [37], PyTorch Geometric [13], and PyTorch Lightning [11] in our implementation.

3.3.4 Dimensionality reduction. PyKale has built a Python version of the MPCA algorithm [35] at `kale.embed.mPCA`, as well as an MPCA-based pipeline at `kale.pipeline.mPCA_trainer`, using both the scikit-Learn library [44] and the TensorLy library [24]. This pipeline has been successfully used for interpretable prediction in gait recognition from video sequences [35], cardiovascular disease diagnosis [55] and prognosis [58] using cardiac magnetic resonance imaging (MRI), and brain state classification using functional MRI (fMRI). We are further building into PyKale other advanced tensor-based algorithms such as regularized Multilinear Regression and Selection (Remurs) [53] and sparse tubal-regularized multilinear regression (Sturm) [29].

3.4 Software engineering

The PyKale team includes machine learning researchers and Research Software Engineers (RSEs). We have adopted good software engineering practices in a research context, often based on other libraries, particularly those in the PyTorch ecosystem.

3.4.1 Version control and collaboration. We use git for version control and GitHub for collaborative working. The PyKale repository stipulates a license (MIT) and contributing guidelines.⁴ These enable reuse of the software and sustainability of the project through

community contributions. This is a platform for long term availability of the resource and lays the foundation for community maintenance over an indefinite period.

3.4.2 Documentation. We use “docstrings” to embed documentation within the source code to maximize synchronicity between code and documentation. Sphinx⁵ is used to automatically build these (along with additional information in “reStructuredText” format) into html docs. Documentation is published via readthedocs.com and is kept up to date by Continuous Integration (CI). This is useful for keeping users and developers up-to-date with new features and bug fixes in a sustainable way. Detailed installation instructions are included.⁶

3.4.3 Tests and continuous integration. We use the PyTest framework and currently have 88% test coverage. The test suite can be run locally, and also runs automatically on GitHub and must pass for code to be merged into the main branch. This ensures that new features do not create unintended side-effects. CI is implemented using GitHub workflows/actions.⁷ Our CI checks include static analysis, pre-commit checks (e.g., maximum file size), documentation building (via Read the Docs), changelog update, project assignment for issues and pull requests, PyPI release of packages, PyTest tests on multiple platforms and multiple python versions, and Codecov code coverage report. This ensures that the version in the main branch is always the most up to date working version, and meets our standards of functionality and coding style.

To maintain a small repository size (currently less than 1MB), we store test data in a separate repository at <https://github.com/pykale/data>, which can be downloaded via `download_file_by_url` in `kale.utils.download` automatically.

4 PYKALE USAGE

Interdisciplinary research is a complex subject to support and care has to be taken to lower the barriers to entry. PyKale includes examples and tutorials to help users' exploration.

4.1 Usage of pipeline-based API in examples

PyKale examples are highly standardized. Each example typically has three essential modules (`main.py`, `config.py`, `model.py`), one optional directory (`configs`), and possibly other modules (`trainer.py`):

- `main.py` is the main module to be run, showing the main workflow.
- `config.py` is the configuration module that sets up the data, prediction problem, hyper-parameters, etc. The settings in this module are the default configuration.
- `configs` is the directory to place customized configurations for individual runs. We use `.yaml` files (see Section 4.3) for this purpose.
- `model.py` is the model module to define the machine learning model and configure its training parameters.
- `trainer.py` is the trainer module to define the training and testing workflow. This module is only needed when NOT using PyTorch Lightning.

⁵<https://www.sphinx-doc.org/>

⁶<https://pykale.readthedocs.io/en/latest/installation.html>

⁷<https://github.com/pykale/pykale/tree/main/.github/workflows>

⁴<https://github.com/pykale/pykale/blob/main/.github/CONTRIBUTING.md>

4.2 Building new modules or projects

Users can build new modules or projects following the steps below.

- Step 1 - Examples: Choose one of the examples of the users' interest (e.g., most relevant to the users' project) to
 - browse through the configuration, main, and model modules,
 - download the data if needed, and
 - run the example following instructions in the example's README.
- Step 2a - New model: To develop new machine learning models under PyKale,
 - define the blocks in the users' pipeline to figure out what the methods are for data loading, preprocessing data, embedding (encoder/representation), prediction (decoder), evaluation, and interpretation, and
 - modify existing pipelines with the users' customized blocks or build a new pipeline with pykale blocks and blocks from other libraries.
- Step 2b - New applications: To develop new applications using PyKale,
 - clarify the input data and the prediction target to find matching functionalities in pykale (request if not found), and
 - tailor data loading, preprocessing, and evaluation (and interpretation if needed) to the users' application.

4.3 YAML configuration

PyKale examples configure a machine learning system using YAML [4]. This is inspired by the usage of YAML in the GitHub package for the Isometric Network (ISONet) [47],⁸ with our adapted version illustrated in Code 2. As modern machine learning systems typically have many settings to configure, specifying many/all settings in command line or Python modules becomes difficult to manage and read. Using YAML greatly improves the readability and reproducibility, and makes configuration changes much easier, via a default configuration specified in `config.py` (top of Code 2) and customized configurations specified in a respective `.yaml` file (bottom of Code 2), which will be merged to overwrite the default setting at run time.

4.4 Notebook tutorials with Binder and Colab

We have eight real-world examples of PyKale usage.⁹ However, tutorials without the need of any installation are important for new users to get familiar with the PyKale workflow and API. For these we must scale-back on real-world datasets due to the computational resources needed, as these lead to long runtimes, unsuitable for interactive learning. Therefore, we are simplifying our examples into Jupyter notebook tutorials so that each tutorial takes minutes instead of hours to run. This will strike a balance between computational requirements and runtime, without resorting to toy examples.

⁸<https://github.com/HaozhiQi/ISONet>

⁹<https://github.com/pykale/pykale/tree/main/examples>

```

1 # The file config.py that defines the default configuration
2 from yacs.config import CfgNode as CN
3
4 # Config definition
5 _C = CN()
6
7 # Dataset
8 _C.DATASET = CN()
9 _C.DATASET.ROOT = "../data"
10 _C.DATASET.NAME = "CIFAR10"
11
12 # Solver
13 _C.SOLVER = CN()
14 _C.SOLVER.SEED = 2020
15 _C.SOLVER.BASE_LR = 0.05
16 _C.SOLVER.TRAIN_BATCH_SIZE = 128
17 _C.SOLVER.MAX_EPOCHS = 100
18
19 # ISONet configs
20 _C.ISON = CN()
21 _C.ISON.DEPTH = 34
22
23 # Misc options
24 _C.OUTPUT_DIR = "../outputs"

```

```

1 # Customization in a .yaml file
2 SOLVER:
3   BASE_LR: 0.01
4   MAX_EPOCHS: 10 # For quick testing
5 ISON:
6   DEPTH: 38

```

Code 2: Code snippets to demonstrate the usage of YAML to configure machine learning systems in PyKale, from `pykale/examples/cifar_isonet/config.py`, which is adapted from <https://github.com/HaozhiQi/ISONet>.

To bring further convenience, we set up cloud-based services with both Binder¹⁰ and Google Colaboratory (Colab)¹¹ for our notebook tutorials so that any users can run PyKale tutorials without the need of any installation. The first such tutorial has been released,¹² with screenshots in Figure 3. More such tutorials are in development.

5 USE CASES: PYKALE EXAMPLES

PyKale currently has example applications from three areas below:

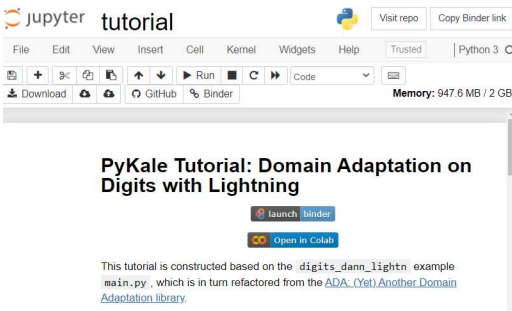
- Image/video recognition: classification of images (objects, digits) or videos (actions in first-person videos);
- Bioinformatics/graph analysis: prediction of links between entities in knowledge graphs (BindingDB, BioSNAP-Decagon);
- Medical imaging: disease diagnosis from cardiac MRIs.

The above examples deal with graphs, images, and videos. Our current APIs support text processing (e.g., for NLP tasks) but an example in this area is still in development. We are also conducting research in integrating audio features in action recognition so an

¹⁰<https://mybinder.org/>

¹¹<https://colab.research.google.com/>

¹²https://github.com/pykale/pykale/blob/main/examples/digits_dann_lightn/tutorial.ipynb



(a) Binder



(b) Google Colab

Figure 3: PyKale digits domain adaptation example on cloud-based services.

example involving audio data is a future task as well. Examples in computer vision applications such as image and video recognition are a good start for most users due to the popularity of vision applications and a low barrier to entry (e.g., no need for specific domain knowledge as in drug discovery). Models first developed in computer vision can be reused or recycled for other applications.

The data used in PyKale examples are real-world data frequently used in research papers. Subsequently, it may take quite some time to finish running these examples. For quick running and demonstration of the workflow, tutorials (Section 4.4) are simplified examples that serve as a better starting point. The following subsections give an overview of the data and algorithms used in the example machine learning systems in PyKale.

5.1 CIFAR and digits classification

Small-image datasets are good for building examples of real-world relevance. PyKale has three such examples: two on the CIFAR datasets [25] and one on digits datasets including MNIST [9, 28], modified MNIST [14], USPS [20], and SVHN [40]. `cifar_isonet` is the first example in PyKale refactoring the ISONet code <https://github.com/HaozhiQi/ISONet> on CIFAR10 and CIFAR100 into the PyKale API. `cifar_cnntransformer` is another example on CIFAR showing the simple application of a mixed CNN and transformer model for vision tasks. Code 1 has shown code snippets of digits classification via domain adaptation using MNIST as the source domain and USPS as the target domain, in the `digits_dann_lightn` example.

Examples on these popular datasets could help users familiar with them make easy connections between the PyKale API and what they are already familiar with.

```

1 # Transform and dataset for multi-modal video data
2 from kale.prepdata.video_transform import get_transform
3 from kale.loaddata.video_datasets import BasicVideoDataset
4
5 transform = get_transform(transform_kind, image_modality)
6 dataset = BasicVideoDataset(data_path, train_list,
7                             imagefile_template="frame_{:010d}.jpg"
8                             if image_modality in ["rgb"]
9                             else "flow_{:010d}.jpg",
10                            transform=transform,
11                            image_modality=image_modality)
12
13 # Action video feature extractor
14 from kale.embed.video_feature_extractor import
15                                     get_video_feat_extractor
16 feature_extractor = get_video_feat_extractor("I3D", image_modality,
17                                             attention, num_classes)

```

Code 3: Code snippets to demonstrate the example on action recognition via domain adaptation.

5.2 Action recognition via domain adaptation

5.2.1 Data. For this action recognition example, we constructed three first-person vision datasets, `ADLsmall`, `GTEA-KITCHEN` and `EPICcopr20`, with 222/454/10094 action videos respectively. We selected and reorganized three videos from ADL dataset [46] as three domains of `ADLsmall`, re-annotated `GTEA` [12, 30] and `KITCHEN` [8] datasets to build `GTEA-KITCHEN`, and adopted the public dataset `EPICcopr20` [39]. We will provide instructions on how to construct these datasets from the source at https://github.com/pykale/data/tree/main/video_data/video_test_data (in progress). Each dataset has two modalities: RGB and optical flow.

5.2.2 Algorithms. For video feature extraction, we built two state-of-the-art action recognition algorithms, `I3D` [7] and `3D ResNet` [57], into PyKale. For domain adaptation, we extended the domain adaptation framework for images (digits, adapted from [56]) to videos. We followed the same pipeline as digits classification while providing additional specific functions for action videos and multi-modal data. As shown in Code 3, the image modality parameter can be set to choose the proper data transform and loader for different modalities: RGB, optical flow, and joint, and all video feature extractors are accessed via a unified interface.

5.3 Drug-target interaction prediction

5.3.1 Data. Predicting the binding affinity between drug compounds and target proteins is fundamental for drug discovery and drug repurposing. This example uses three public datasets (for three metrics `Ki`, `Kd` and `IC50`) from `BindingDB` [31] containing 52,284, 375,032, and 991,486 interaction pairs respectively, accessed via the `Therapeutics Data Commons (TDC)` platform [19]. Given the amino acid sequences of targets and `SMILES` (Simplified Molecular Input Line Entry System) strings of drug compounds, the task is to predict drug-target binding affinity. Following [22], the affinity metrics are transformed into the logarithm form for more stable training and validation.

```

1 from kale.loaddata.tdc_datasets import BindingDBDataset
2 from kale.embed.seq_nn import CNNEncoder
3 from kale.predict.decode import MLPDecoder
4 from kale.pipeline.deep_dti import DeepDTATrainer
5 import pytorch_lightning as pl
6
7 # Load training, validation, and test data
8 train_dataset = BindingDBDataset(name="Kd", split="train")
9 val_dataset = BindingDBDataset(name="Kd", split="valid")
10 test_dataset = BindingDBDataset(name="Kd", split="valid")
11 train_loader = DataLoader(dataset=train_dataset, batch_size=64)
12 val_loader = DataLoader(dataset=val_dataset, batch_size=64)
13 test_loader = DataLoader(dataset=test_dataset, batch_size=64)
14
15 # Initialize DeepDTA with encoder and decoder
16 drug_encoder, target_encoder = CNNEncoder(), CNNEncoder()
17 decoder = MLPDecoder()
18 model = DeepDTATrainer(drug_encoder, target_encoder, decoder,
19                       ci_metric=True)
20
21 # Train the model under the PyTorch Lightning framework
22 trainer = pl.Trainer(max_epochs=100)
23 trainer.fit(model, train_dataloader=train_loader,
24            val_data_loaders=val_loader)
25 trainer.test(test_data_loaders=test_loader)

```

Code 4: Demonstration code for drug-target interaction prediction with DeepDTA.

5.3.2 Algorithms. For this drug-target interaction prediction problem, we built DeepDTA [42] into PyKale, a typical CNN-based model with encoder-decoder architecture. The drug SMILES string and target amino acid sequence are encoded by their independent CNNs, and then a multilayer perceptron is used to decode the affinity from the drug-target combined encoding. We refactored DeepDTA into the PyKale API structure with separate modules for load data, preprocess data, encode (embed) and decode (predict). These refactored modules are flexible and can be reused across different applications. Code 4 illustrates the usage of the DeepDTA model as implemented in PyKale.

5.4 Polypharmacy side effect prediction via knowledge graph link prediction

5.4.1 Data. Polypharmacy uses drug combination to treat complex diseases, which may cause side effects. Predicting such side effects can be formulated as a link prediction problem on knowledge graphs of drugs and proteins [41, 61, 62]. This example uses the public BioSNAP-Decagon dataset [63] with 6,075,428 edges, 1,100 different edge labels, and three types of edges: drug-drug interaction, protein-protein interaction and drug-target interaction. The drug-drug interactions model polypharmacy side effects.

5.4.2 Algorithms. We built our recent GripNet [61] into PyKale for this example. GripNet is a subgraph network framework for multi-relational link prediction on heterogeneous graphs via segregated node representation learning on “supervertices” and “superedges”. PyKale implements APIs to support these advanced concepts for node embedding and link prediction on heterogeneous graphs.

```

1 from kale.pipeline.mpca_trainer import MPCATrainer
2 from kale.interpret import model_weights
3 from sklearn.model_selection import train_test_split
4 from sklearn.metrics import accuracy_score, roc_auc_score
5
6 # Assume data x and labels y are loaded and preprocessed
7 trainer = MPCATrainer(mpca_params={"return_vector": True})
8 x_train, x_test, y_train, y_test = train_test_split(x, y,
9                                                    test_size=0.2)
10 trainer.fit(x_train, y_train)
11 print("Accuracy:", accuracy_score(y_test, trainer.predict(x_test)))
12 print("AUC:", roc_auc_score(y_test,
13                             trainer.decision_function(x_test)))
14
15 # Visualize model coefficients (weights) for interpretation
16 weights = trainer.mpca.inverse_transform(trainer.clf.coef_)
17         - trainer.mpca.mean_
18 top_weights = model_weights.select_top_weight(weights)
19 model_weights.plot_weights(top_weights[0][0],
20                            background_img=x[0][0])

```

Code 5: MPCA-based pipeline for predicting pulmonary arterial hypertension with cardiac MRI.

5.5 Cardiac MRI classification

5.5.1 Data. The data used to build this example are a subset of the dataset used in [55], which consists of cardiac MRI (CMRI) sequences acquired from patients with pulmonary arterial hypertension and health controls. They are not yet in the public domain but we are exploring release options. The CMRI sequences are standardized using methods in `kale.prepdata.image_transform`.

5.5.2 Algorithms. This example uses the MPCATrainer pipeline in `kale.pipeline.mpca_trainer`, which implemented the machine learning pipeline used in [35, 54, 55] in the scikit-learn style. The three key steps are MPCA dimensionality reduction, feature selection, and classification (SVM, Linear SVM, or logistic regression), where the feature selection and classification algorithms reuse respective APIs in scikit-learn [44]. Code 5 shows the steps for MPCA-based prediction and interpretation on cardiac MRI.

6 PYKALE OPENNESS AND PLAN

6.1 License and community engagement

PyKale is publicly available at <https://github.com/pykale/pykale> under an MIT license, which is a simple permissive license with minimal restrictions. The PyKale GitHub repository has active discussion board at <https://github.com/pykale/pykale/discussions> and project board at <https://github.com/pykale/pykale/projects> to interact with users and make the development process and plan transparent to all users. Complete documentation is hosted at <https://pykale.readthedocs.io/> with multiple versions available, generated automatically from <https://github.com/pykale/pykale/tree/main/docs>. We have provided tutorials and examples as well as detailed contributing guidelines at <https://github.com/pykale/pykale/blob/main/.github/CONTRIBUTING.md> and change logs at <https://github.com/pykale/pykale/blob/main/.github/CHANGELOG.md>.

We also released a 12-minute YouTube video at <https://youtu.be/i5BYdMfbpMQ> to briefly explain the motivation and principles

behind PyKale. This paper is another effort to reach out to the wider research community to share this resource and get feedback for further improvements.

6.2 Limitations and future development

PyKale is an open-source project started in June 2020, with the first PyPI release in January 2021. It was motivated by the growing needs for machine learning systems that can deal with multiple sources of data, particularly in interdisciplinary areas such as healthcare. For example, clinicians often need to make use of a combination of medical images (e.g., X-rays, CTs, MRIs), biological data (gene, DNA, RNA), and electronic health record for decision making. Our focus on multimodal learning and transfer learning has defined a challenging scope, while holding the promises to break barriers in interdisciplinary research.

To date, PyKale has built APIs supporting machine learning from graphs, images, texts, and videos, with four mature pipelines implemented. Nevertheless, we do not have an example on text data yet, and we have not built APIs for audio yet. Developing projects involving multiple data sources takes considerably longer time than developing those involving a single data source. The current version of PyKale has two examples on multimodal learning involving heterogeneous drug and target data and two examples on domain adaptation (transfer learning) for images and videos. These examples laid solid foundations for us to grow our research in these areas and build more advanced examples in future development. In addition, our tests currently have a coverage of 88%. We need further improvements for a higher coverage and more rigorous tests.

7 CONCLUSIONS

In this paper, we have introduced *PyKale*, a Python library for knowledge-aware machine learning from multiple sources, particularly from multiple modalities for multimodal learning and from multiple domains for transfer learning. This library was motivated by needs in healthcare applications (hence the acronym *kale*, a healthy vegetable) and aims to enable and accelerate interdisciplinary research. Building on standard software engineering practices, we proposed a new green machine learning perspective to advocate reducing repetitions and redundancy, reusing existing resources, and recycling learning models across areas. Following such principles, we designed our API to be pipeline-based to unify the workflow and increase the flexibility. This design can help to break barriers between different areas or applications and facilitate the fusion and nurture of ideas across discipline boundaries.

The goal of PyKale is to facilitate interdisciplinary, knowledge-aware machine learning research for graphs, images, texts, and videos. It will make it easier to bring machine learning models developed in one area to the other, and integrate data from multiple sources for prediction tasks in interdisciplinary areas. Its focus on leveraging knowledge from multiple sources also helps accurate and interpretable prediction. To demonstrate such potential, we have shown example applications including bioinformatics, knowledge graph, image/video recognition, and medical imaging on real-world datasets.

ACKNOWLEDGMENTS

The development of PyKale is partially supported by the Innovator Awards: Digital Technologies from the Wellcome Trust (grant 215799/Z/19/Z). We thank the support and contributions from David Jones, Will Furnass, and other members of the Research Software Engineering (RSE) team headed by Paul Richmond. We also thank early users of our library and their helpful feedback.

REFERENCES

- [1] Kartik Ahuja and Mihaela van der Schaar. 2019. Joint Concordance Index. In *Proceedings of the 2019 53rd Asilomar Conference on Signals, Systems, and Computers*. 2206–2213.
- [2] Peizhen Bai, Yan Ge, Fangling Liu, and Haiping Lu. 2019. Joint interaction with context operation for collaborative filtering. *Pattern Recognition* 88 (2019), 729–738.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Fernando Pereira, et al. 2007. Analysis of representations for domain adaptation. In *Proceedings of the Advances in Neural Information Processing Systems*. 137–144.
- [4] Oren Ben-Kiki, Clark Evans, and Brian Ingerson. 2009. Yaml ain't markup language (yaml™) version 1.1. *Working Draft 2008-05 11* (2009).
- [5] Antonio Candelieri, Riccardo Perego, and Francesco Archetti. 2021. Green machine learning via augmented Gaussian processes and multi-information source optimization. *Soft Computing* (2021), 1–13.
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*. 213–229.
- [7] Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6299–6308.
- [8] Fernando De la Torre, Jessica Hodgins, Adam Bargteil, Xavier Martin, Justin Macey, Alex Collado, and Pep Beltran. 2009. Guide to the carnegie mellon university multimodal activity (cmu-mmact) database. (2009).
- [9] Li Deng. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 6 (2012), 141–142.
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, and et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proceedings of the International Conference on Learning Representations*.
- [11] William A Falcon and et al. 2019. PyTorch Lightning. *GitHub*. 3 (2019). <https://github.com/PyTorchLightning/pytorch-lightning>
- [12] Alireza Fathi, Xiaofeng Ren, and James M Rehg. 2011. Learning to recognize objects in egocentric activities. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3281–3288.
- [13] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [14] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17, 1 (2016), 2096–2030.
- [15] Eva Garcia Martin. 2017. Energy efficiency in machine learning: A position paper. In *Proceedings of the 30th Annual Workshop of the Swedish Artificial Intelligence Society*, Vol. 137. 68–72.
- [16] Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. 2018. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. In *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 31. 7587–7597.
- [17] Georgian. 2020. Multimodal-Toolkit. *GitHub* (2020). <https://github.com/georgiano/Multimodal-Toolkit>
- [18] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7132–7141.
- [19] Kexin Huang, Tianfan Fu, Wenhao Gao, and et al. 2021. Therapeutics Data Commons: Machine Learning Datasets and Tasks for Therapeutics. *arXiv preprint arXiv:2102.09548* (2021).
- [20] Jonathan Hull. 1994. A database for handwritten text recognition research. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 16, 5 (1994), 550–554.
- [21] Junguang Jiang, Bo Fu, and Mingsheng Long. 2020. Transfer-Learning-Library. *GitHub* (2020). <https://github.com/thuml/Transfer-Learning-Library>
- [22] Mostafa Karimi, Di Wu, Zhangyang Wang, and Yang Shen. 2019. DeepAffinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics* 35, 18 (2019), 3329–3338.
- [23] Thomas Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the 5th International Conference on Learning Representations*.

- [24] Jean Kossaifi, Yannis Panagakis, Anima Anandkumar, and Maja Pantic. 2019. TensorLy: Tensor Learning in Python. *Journal of Machine Learning Research* 20, 26 (2019), 1–6.
- [25] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. 2010. Cifar-10 (canadian institute for advanced research). 5 (2010).
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Proceedings of the Advances in Neural Information Processing Systems*. 1097–1105.
- [27] Sun Yuan Kung. 2014. *Kernel methods and machine learning*. Cambridge University Press.
- [28] Yann LeCun. 1998. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998).
- [29] Wenwen Li, Jian Lou, Shuo Zhou, and Haiping Lu. 2019. Sturm: Sparse tubular-regularized multilinear regression for fmri. In *Proceedings of the International Workshop on Machine Learning in Medical Imaging*. 256–264.
- [30] Yin Li, Zhefan Ye, and James M Rehg. 2015. Delving into egocentric actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 287–295.
- [31] Tiqing Liu, Yuhmei Lin, Xin Wen, R. Jorissen, and M. Gilson. 2007. BindingDB: a web-accessible database of experimentally determined protein–ligand binding affinities. *Nucleic Acids Research* 35 (2007), D198 – D201.
- [32] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. 2015. Learning transferable features with deep adaptation networks. In *Proceedings of the International Conference on Machine Learning*. 97–105.
- [33] Mingsheng Long, ZHANGJIE CAO, Jianmin Wang, and Michael I Jordan. 2018. Conditional Adversarial Domain Adaptation. In *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 31.
- [34] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. 2017. Deep transfer learning with joint adaptation networks. In *Proceedings of the International Conference on Machine Learning*. 2208–2217.
- [35] Haiping Lu, Konstantinos N Plataniotis, and Anastasios N Venetsanopoulos. 2008. MPCA: Multilinear principal component analysis of tensor objects. *IEEE Transactions on Neural Networks* 19, 1 (2008), 18–39.
- [36] Nic Ma, Wenqi Li, and Richard Brown. 2021. *Project-MONAI/MONAI: 0.5.3*. <https://doi.org/10.5281/zenodo.4891800>
- [37] Sébastien Marcel and Yann Rodriguez. 2010. Torchvision the Machine-Vision Package of Torch. In *Proceedings of the 18th ACM International Conference on Multimedia*. 1485–1488.
- [38] Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, et al. 2016. Mllib: Machine learning in apache spark. *Journal of Machine Learning Research* 17, 1 (2016), 1235–1241.
- [39] Jonathan Munro and Dima Damen. 2020. Multi-modal domain adaptation for fine-grained action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 122–132.
- [40] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning. *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning* (2011).
- [41] Vit Nováček and Sameh K Mohamed. 2020. Predicting Polypharmacy Side-effects Using Knowledge Graph Embeddings.. In *Proceedings of the AMIA Joint Summits on Translational Science*. 449–458.
- [42] Hakime Öztürk, E. Olmez, and Arzucan Özgür. 2018. DeepDTA: deep drug–target binding affinity prediction. *Bioinformatics* 34 (2018), i821 – i829.
- [43] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. 2010. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* 22, 2 (2010), 199–210.
- [44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [45] Fernando Pérez-García, Rachel Sparks, and Sebastien Ourselin. 2020. TorchIO: a Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. (2020). <http://arxiv.org/abs/2003.04696>
- [46] Hamed Pirsiavash and Deva Ramanan. 2012. Detecting activities of daily living in first-person camera views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2847–2854.
- [47] Haozhi Qi, Chong You, Xiaolong Wang, Yi Ma, and Jitendra Malik. 2020. Deep isometric learning for visual recognition. In *Proceedings of the International Conference on Machine Learning*. 7824–7835.
- [48] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. 2020. Kornia: an open source differentiable computer vision library for pytorch. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*. 3674–3683.
- [49] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. 234–241.
- [50] Aghiles Salah, Quoc-Tuan Truong, and Hady W Lauw. 2020. Cornac: A Comparative Framework for Multimodal Recommender Systems. *Journal of Machine Learning Research* 21, 95 (2020), 1–5.
- [51] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. 2018. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [52] Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2020. MMF: A multimodal framework for vision and language research. <https://github.com/facebookresearch/mmf>.
- [53] Xiaonan Song and Haiping Lu. 2017. Multilinear regression for embedded feature selection with application to fmri analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [54] Xiaonan Song, Lingnan Meng, Qiquan Shi, and Haiping Lu. 2015. Learning tensor-based features for whole-brain fMRI classification. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. 613–620.
- [55] Andrew J Swift, Haiping Lu, Johanna Uthoff, Pankaj Garg, Marcella Coglianò, Jonathan Taylor, Peter Metherall, Shuo Zhou, Christopher S Johns, Samer Alabed, et al. 2021. A machine learning cardiac magnetic resonance approach to extract disease features and automate pulmonary arterial hypertension diagnosis. *European Heart Journal-Cardiovascular Imaging* 22, 2 (2021), 236–245.
- [56] Anne-Marie Tusch and Christophe Renaudin. 2020. (Yet) Another Domain Adaptation library. <https://github.com/criteo-research/pytorch-ada>
- [57] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. 2018. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6450–6459.
- [58] Johanna Uthoff, Samer Alabed, Andrew J Swift, and Haiping Lu. 2020. Geodesically Smoothed Tensor Features for Pulmonary Hypertension Prognosis Using the Heart and Surrounding Tissues. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*. 253–262.
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan Gomez, Ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the Advances in Neural Information Processing Systems*. 6000–6010.
- [60] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. 2018. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7794–7803.
- [61] Hao Xu, Shengqi Sang, Peizhen Bai, Laurence Yang, and Haiping Lu. 2020. GripNet: Graph Information Propagation on Supergraph for Heterogeneous Graphs. arXiv:2010.15914 [cs.LG]
- [62] M. Zitnik, Monica Agrawal, and J. Leskovec. 2018. Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics* 34 (2018), i457 – i466.
- [63] Marinka Zitnik, Rok Sosič, Sagar Maheshwari, and Jure Leskovec. 2018. BioSNAP Datasets: Stanford Biomedical Network Dataset Collection.
- [64] Hakime Öztürk, Arzucan Özgür, and Elif Ozkirimli. 2018. DeepDTA: deep drug–target binding affinity prediction. *Bioinformatics* 34, 17 (2018), i821–i829.