

# Latte: cross-framework Python package for evaluation of latent-based generative models

Watcharasupat, Karn N.; Lee, Junyoung; Lerch, Alexander

2022

Watcharasupat, K. N., Lee, J. & Lerch, A. (2022). Latte: cross-framework Python package for evaluation of latent-based generative models. *Software Impacts*, 11, 100222.  
<https://dx.doi.org/10.1016/j.simpa.2022.100222>

<https://hdl.handle.net/10356/164060>

<https://doi.org/10.1016/j.simpa.2022.100222>

---

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

*Downloaded on 09 Apr 2024 11:37:48 SGT*



Original software publication

## Latte: Cross-framework Python package for evaluation of latent-based generative models

Karn N. Watcharasupat <sup>a,b,\*</sup>, Junyoung Lee <sup>b</sup>, Alexander Lerch <sup>a</sup><sup>a</sup> Center for Music Technology, Georgia Institute of Technology, Atlanta, GA, USA<sup>b</sup> School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore

## ARTICLE INFO

## Keywords:

Deep generative networks  
Disentanglement learning  
Latent space  
Controllable generation  
Python

## ABSTRACT

*Latte* (for *LATent Tensor Evaluation*) is a Python library for evaluation of latent-based generative models in the fields of disentanglement learning and controllable generation. *Latte* is compatible with both PyTorch and TensorFlow/Keras, and provides both functional and modular APIs that can be easily extended to support other deep learning frameworks. Using NumPy-based and framework-agnostic implementation, *Latte* ensures reproducible, consistent, and deterministic metric calculations regardless of the deep learning framework of choice.

## Code metadata

Current code version	v0.1.0
Permanent link to code/repository used for this code version	<a href="https://github.com/SoftwareImpacts/SIMPAC-2021-192">https://github.com/SoftwareImpacts/SIMPAC-2021-192</a>
Permanent link to Reproducible Capsule	<a href="https://codeocean.com/capsule/3186531/tree/v1">https://codeocean.com/capsule/3186531/tree/v1</a>
Legal Code License	MIT License
Code versioning system used	Git
Software code languages, tools, and services used	Language: Python 3 CI/CD: pytest, CircleCI, CodeCov, CodeFactor Python ≥ 3.7, < 3.10 NumPy ≥ 1.18, Scikit-learn ≥ 1.0.0 (optional) PyTorch ≥ 1.3.1, TorchMetrics ≥ 0.2.0 (optional) TensorFlow ≥ 2.0
Compilation requirements, operating environments & dependencies	<a href="https://latte.readthedocs.io/">https://latte.readthedocs.io/</a> <a href="https://github.com/karnwatcharasupat/latte/issues">https://github.com/karnwatcharasupat/latte/issues</a>
If available Link to developer documentation/manual	
Support email for questions	

## 1. Introduction

Disentanglement learning and controllable generation are fast-growing fields within deep learning research, powered by the advances in deep generative networks, such as variational autoencoders (VAEs) [1] and generative adversarial networks (GANs) [2]. Disentanglement learning is often used with encoder-decoder architectures to produce latent representations in the form of latent vectors or tensors in the bottleneck layer, such that each latent dimension has an approximately exclusive mapping to a semantic attribute of interest.

These disentangled latent representations are particularly useful in the generative models that aim to produce samples with specific and controllable semantic attributes [3,4].

With the growth of the fields comes the need for a reliable and consistent method of evaluation that allows for the comparison of different systems across a variety of metrics. Therefore, we introduce *Latte*<sup>1</sup> (for *LATent Tensor Evaluation*), a cross-framework Python package for evaluation of latent-based generative models. Since successful latent-based controllable generation requires more than disentanglement [5], *Latte*

The code (and data) in this article has been certified as Reproducible by Code Ocean: (<https://codeocean.com/>). More information on the Reproducibility Badge Initiative is available at <https://www.elsevier.com/physical-sciences-and-engineering/computer-science/journals>.

\* Corresponding author at: School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore.

E-mail addresses: [kwatcharasupat@gatech.edu](mailto:kwatcharasupat@gatech.edu) (K.N. Watcharasupat), [junyoung002@e.ntu.edu.sg](mailto:junyoung002@e.ntu.edu.sg) (J. Lee), [alexander.lerch@gatech.edu](mailto:alexander.lerch@gatech.edu) (A. Lerch).

<sup>1</sup> Software DOI: [10.5281/zenodo.5786402](https://doi.org/10.5281/zenodo.5786402)

<https://doi.org/10.1016/j.simpa.2022.100222>

Received 20 December 2021; Received in revised form 3 January 2022; Accepted 4 January 2022

also covers interpolability metrics, in addition to disentanglement metrics.

Framework-agnostic evaluation tools are known to greatly facilitate and accelerate research development in a field. For example, in the field of audio source separation, `bss_eval` [6] and its successor `museval` [7] have greatly benefited the field and provided a standard benchmarking tool. Many studies on disentanglement learning have formally or informally relied on the `disentanglement_lib`<sup>2</sup> library for their evaluation. However, the library has not had a new release since 2019. Moreover, `disentanglement_lib` was mainly created as a code base for reproducing the studies [8–13], thus the metric implementations were written to fit the systems under development rather than to cater to a wider range of applications, and are only available in TensorFlow. As a result, researchers working with PyTorch or other incompatible models often have to rely on their own re-implementations of the metrics — an error-prone and inefficient approach that in the best case produces additional work, and in the worst case leads to inconsistencies in evaluation metric implementations. In addition, to the best of our knowledge, no comprehensive library for the evaluation of generative interpolability currently exists.

The introduction of *Latte* aims to address these shortcomings. By design, *Latte* performs all metric calculations with NumPy-based computation to ensure cross-framework consistency. The modular design used in *Latte* also ensures easy extensibility for supporting other frameworks beyond PyTorch and TensorFlow in the future.

## 2. Software description

*Latte* is a cross-framework Python package for the evaluation of disentanglement and controllability in latent-based generative models. *Latte* supports on-the-fly metric calculation for disentanglement learning and controllable generation using both a standalone functional API, and modular APIs for the two major deep learning frameworks — TensorFlow/Keras [14] and PyTorch [15].

In order to maximize cross-framework compatibility and reproducibility, core functionalities of *Latte* are developed with NumPy [16] and Scikit-learn [17], without any deep learning dependencies. These NumPy-based functionalities also serve as a standalone functional API, allowing the use of *Latte* in post-hoc analyses without the need for specific deep learning dependencies like TensorFlow or PyTorch.

For use with deep learning frameworks, we implemented a modular API for the metrics to allow easy usage within the respective framework. For use with TensorFlow/Keras, we implemented wrappers based on the Keras Metric API to convert the core NumPy-based functionalities to TensorFlow-compatible operations. With PyTorch, we implemented similar wrappers based on the TorchMetrics API [18], which allows easy integration with both PyTorch and the popular PyTorch Lightning [19] frameworks. *Latte* modular metrics can be used in distributed training using the respective built-in multi-node support in Keras and TorchMetrics. An example of using *Latte* in modular mode with PyTorch is shown in Fig. 1.

### 2.1. Deterministic metric calculation

A number of metrics used in disentanglement learning and controllable generation are based on randomly-initialized regressors and classifiers [20]. Moreover, practical calculation of probabilistic measures via Scikit-learn, such as mutual information and entropy, also requires random number generation. As `disentanglement_lib` does not explicitly set a seed before metric calculation, identical inputs may result in different metric values. This particular detail is not commonly known amongst end-users who may not be aware of the implementation details. In *Latte*, a random seed of 42 is set by default, but can be switched off by calling `latte.seed(None)`. This allows end-users to have deterministic metric calculation by default without having to know the implementation details of *Latte* and its dependencies.

```
import latte
from latte.metrics.torch.disentanglement
    import MutualInformationGap
import torch

data = initialize_dataset()
model = initialize_model()

mig = MutualInformationGap(
    reg_dim=[0, 3, 7],
    discrete=False
)

for x, attributes in range(data):
    xhat, z = model(x)
    mig.update(z=z, a=attributes)

mig_values = mig.compute()
```

**Fig. 1.** Example of using *Latte* in modular mode with PyTorch via the TorchMetrics API. In this example, the data contains three continuous semantic attributes, which are respectively regularized by the latent specified by the argument `reg_dim`. The `discrete=False` option specifies that the semantic attributes are continuous-valued.

### 2.2. Metric bundles

In addition to individual metric functions and modules, *Latte* provides metric bundles, which are special modules containing multiple metrics commonly used together, similar to `MetricCollection` in TorchMetrics. All metric submodules of a metric bundle are initialized together with a common set of settings, ensuring consistency and compatibility between the metrics within a bundle. Inputs of the update calls to a metric bundle are also automatically passed to the respective submodules, reducing the amount of codes needed for metric calculations. Custom bundles can also be created via the `MetricBundle` class in *Latte*. An example of using a *Latte* metric bundle with TensorFlow/Keras is shown in Fig. 2.

### 2.3. Testing and deployment

Automated testing of *Latte* is performed via `pytest`. Continuous integration and deployment (CI/CD) is handled via CircleCI. Code coverage and code quality are respectively monitored via `CodeCov` and `CodeFactor`. *Latte* releases are available on the Python Package Index (PyPI)<sup>3</sup> and can be easily installed via `pip install latte-metrics`.

## 3. Supported metrics

Successful latent-based controllable generation requires more than disentanglement [5]. Evaluation of a controllable generative system generally falls into three categories: reconstruction, disentanglement, and interpolability. *Latte* currently focuses on disentanglement metrics and interpolability metrics, since evaluation of reconstruction fidelity is domain-specific. Domain-specific reconstruction metrics for generative models may be added in future versions.

<sup>2</sup> [https://github.com/google-research/disentanglement\\_lib](https://github.com/google-research/disentanglement_lib)

<sup>3</sup> <https://pypi.org/project/latte-metrics/>

```

import latte
from latte.metrics.keras.bundles import
    DependencyAwareMutualInformationBundle
    as DAMIBundle
import tensorflow as tf

data = initialize_dataset()
model = initialize_model()

bundle = DAMIBundle(
    reg_dim=[0, 3, 7],
    discrete=False
)

for x, attributes in range(data):
    xhat, z = model(x)
    bundle.update_state(z=z, a=attributes)

metrics = bundle.result()

dmig_values = metrics['DMIG']
dlig_values = metrics['DLIG']

```

**Fig. 2.** Example of using a *Latte* metric bundle in modular mode with TensorFlow via the Keras Metric API. The call signature is very similar to single-metric modules — the main difference is that a metric bundle returns a dictionary of arrays instead of a single array. *DependencyAwareMutualInformationBundle* contains MIG, DMIG, XMIG, and DLIG. All individual metric submodules are automatically initialized with the same `reg_dim` and `discrete` options.

As of the current version (v0.1.0), the following disentanglement metrics are supported: Mutual Information Gap (MIG) [21], Separate Attribute Predictability (SAP) [20], Modularity [22], Dependency-aware Mutual Information Gap (DMIG) [23], Dependency-blind Mutual Information Gap (XMIG), and Dependency-aware Latent Information Gap (DLIG) [24]. The interpolability metrics supported are: Smoothness, and Monotonicity [24]. We briefly describe the supported metrics in the following sections.

### 3.1. Disentanglement metrics

**Mutual Information Gap.** Mutual Information Gap (MIG) evaluates the degree of disentanglement of a latent vector  $\mathbf{z} \in \mathbb{R}^D$  with respect to a particular semantic attribute,  $a_i \in \mathbb{R}$ , by considering the gap in mutual information between the attribute and its most informative latent dimension and that between the attribute and its second-most informative latent dimension [21]. Mathematically, MIG is given by

$$\text{MIG}(a_i, \mathbf{z}) = \frac{\mathcal{I}(a_i, z_j) - \mathcal{I}(a_i, z_k)}{\mathcal{H}(a_i)}, \quad (1)$$

where  $j = \text{argmax}_d \mathcal{I}(a_i, z_d)$ ,  $k = \text{argmax}_{d \neq j} \mathcal{I}(a_i, z_d)$ ,  $\mathcal{I}(\cdot, \cdot)$  is mutual information, and  $\mathcal{H}(\cdot)$  is entropy.

**Separate Attribute Predictability.** Separate Attribute Predictability (SAP) is similar in nature to MIG but, instead of mutual information, uses the coefficient of determination for continuous attributes and classification accuracy for discrete attributes to measure the extent of relationship between a latent dimension and an attribute [20]. SAP is given by

$$\text{SAP}(a_i, \mathbf{z}) = S(a_i, z_j) - S(a_i, z_k), \quad (2)$$

where  $j = \text{argmax}_d S(a_i, z_d)$ ,  $k = \text{argmax}_{d \neq j} S(a_i, z_d)$ , and  $S(\cdot, \cdot)$  is either the coefficient of determination or classification accuracy.

**Modularity.** Modularity is a latent-centric measure of disentanglement [22] based on mutual information. Modularity measures the degree in which a latent dimension contains information about only one attribute, and is given by

$$\text{Modularity}(\{a_i\}, z_d) = 1 - \frac{\sum_{i \neq j} (\mathcal{I}(a_i, z_d) / \mathcal{I}(a_j, z_d))^2}{|\{a_i\}| - 1}, \quad (3)$$

where  $j = \text{argmax}_i \mathcal{I}(a_i, z_d)$ .

**Dependency-aware Mutual Information Gap.** Dependency-aware Mutual Information Gap (DMIG) is a dependency-aware version of MIG that accounts for attribute interdependence observed in real-world data [23]. Mathematically, DMIG is given by

$$\text{DMIG}(a_i, \mathbf{z}) = \frac{\mathcal{I}(a_i, z_j) - \mathcal{I}(a_i, z_k)}{\mathcal{H}(a_i | a_l)}, \quad (4)$$

where  $j = \text{argmax}_d \mathcal{I}(a_i, z_d)$ ,  $k = \text{argmax}_{d \neq j} \mathcal{I}(a_i, z_d)$ ,  $\mathcal{H}(\cdot | \cdot)$  is conditional entropy, and  $a_l$  is the attribute regularized by  $z_k$ . If  $z_k$  is not regularizing any attribute, DMIG reduces to the usual MIG. DMIG compensates for the reduced maximum possible value of the numerator due to attribute interdependence.

**Dependency-blind Mutual Information Gap.** Dependency-blind Mutual Information Gap (XMIG) is a complementary metric to MIG and DMIG that measures the gap in mutual information with the subtrahend restricted to dimensions which do not regularize any attribute [24]. XMIG is given by

$$\text{XMIG}(a_i, \mathbf{z}) = \frac{\mathcal{I}(a_i, z_j) - \mathcal{I}(a_i, z_k)}{\mathcal{H}(a_i)}, \quad (5)$$

where  $j = \text{argmax}_d \mathcal{I}(a_i, z_d)$ ,  $k = \text{argmax}_{d \notin D} \mathcal{I}(a_i, z_d)$ , and  $D$  is a set of latent indices which do not regularize any attribute. XMIG allows monitoring of latent disentanglement exclusively against non-regularized latent dimensions.

**Dependency-aware Latent Information Gap.** Dependency-aware Latent Information Gap (DLIG) is a latent-centric counterpart to DMIG [24]. DLIG evaluates disentanglement of a set of semantic attributes  $\{a_i\}$  with respect to a latent dimension  $z_d$ , such that

$$\text{DLIG}(\{a_i\}, z_d) = \frac{\mathcal{I}(a_j, z_d) - \mathcal{I}(a_k, z_d)}{\mathcal{H}(a_j | a_k)}, \quad (6)$$

where  $j = \text{argmax}_i \mathcal{I}(a_i, z_d)$ ,  $k = \text{argmax}_{i \neq j} \mathcal{I}(a_i, z_d)$ .

### 3.2. Interpolability metrics

The two interpolability metrics currently supported are based on the concept of a pseudo-derivative called latent-induced attribute difference (LIAD), which is defined as

$$D_{i,d}(\mathbf{z}; \delta) = \frac{\mathcal{A}_i(\mathbf{z} + \delta e_d) - \mathcal{A}_i(\mathbf{z})}{\delta} \quad (7)$$

where  $\mathcal{A}_i(\cdot)$  is the measurement of attribute  $a_i$  from a sample generated from its latent vector argument,  $d$  is the latent dimension regularizing  $a_i$ ,  $\delta > 0$  is the latent step size, and  $e_d$  is the  $d$ th elementary vector [24]. Second-order LIAD is similarly defined by

$$D_{i,d}^{(2)}(\mathbf{z}; \delta) = \frac{D_i^{(1)}(\mathbf{z} + \delta e_d) - D_i^{(1)}(\mathbf{z})}{\delta}, \quad (8)$$

where  $D^{(1)} \equiv D$ .

**Smoothness.** Smoothness is a measure of how smoothly an attribute changes with respect to a change in the regularizing latent dimension [24]. Smoothness of a latent vector  $\mathbf{z}$  is based on the concept of second-order derivative, and is given by

$$\text{Smoothness}_{i,d}(\mathbf{z}; \delta) = 1 - \frac{\mathcal{C}_{k \in \mathfrak{K}} \left[ D_{i,d}^{(2)}(\mathbf{z} + k\delta e_d; \delta) \right]}{\delta^{-1} \mathcal{R}_{k \in \mathfrak{K}} \left[ D_{i,d}^{(1)}(\mathbf{z} + k\delta e_d; \delta) \right]}, \quad (9)$$

where  $\mathcal{C}_{k \in \mathfrak{K}}[\cdot]$  is the contraharmonic mean of its arguments over values of  $k \in \mathfrak{K}$ , and  $\mathcal{R}_{k \in \mathfrak{K}}[\cdot]$  is the range of its arguments over values of  $k \in \mathfrak{K}$ , and  $\mathfrak{K}$  is the set of interpolating points used during evaluation.

**Monotonicity.** Monotonicity is a measure of how monotonic an attribute changes with respect to a change in the regularizing dimension [24]. Monotonicity of a latent vector  $\mathbf{z}$  is given by

$$\text{Monotonicity}_{i,d}(\mathbf{z}; \delta, \epsilon) = \frac{\sum_{k \in \mathfrak{K}} \mathbb{I} \left[ |D_{i,d}^{(1)}(\mathbf{z} + k\delta e_d; \delta)| > \epsilon \right] \cdot S_k}{\sum_{k \in \mathfrak{K}} \mathbb{I} \left[ |D_{i,d}^{(1)}(\mathbf{z} + k\delta e_d; \delta)| > \epsilon \right]}, \quad (10)$$

where  $S_k = \text{sgn}(D_{i,d}^{(1)}(\mathbf{z} + k\delta e_d; \delta)) \in \{-1, 0, 1\}$ ,  $\mathbb{I}[\cdot]$  is the Iverson bracket operator, and  $\epsilon > 0$  is a noise threshold for ignoring near-zero attribute changes.

#### 4. Software impact and future work

*Latte* is released under the MIT License and welcomes community contribution to the package. The authors hope that the introduction of *Latte* will reduce the amount of time spent on re-implementing evaluation metrics due to framework incompatibility, and provide a standardized and uniform framework for evaluation of controllable generative systems regardless of the deep learning framework of choice.

The implementation of metrics such as Latent Density Ratio [5] and Linearity [4] is currently planned for future releases. Additional metrics under consideration include  $\beta$ -VAE Score [1]; FactorVAE Score [25]; Explicitness [22]; Disentanglement, Completeness and Informativeness (DCI) Scores [26]; Interventional Robustness Score (IRS) [27]; Consistency and Restrictiveness [12]; and, Density and Coverage [28]. Wrapper supports for PyTorch Ignite [29] and nnabla [30] are also currently under consideration.

#### 5. Conclusion

We introduce *Latte*, a cross-framework Python package for evaluation of latent-based generative models. *Latte* supports on-the-fly metric calculation for disentanglement learning and controllable generation using both standalone functional API based on NumPy, and modular APIs for both TensorFlow/Keras and PyTorch. *Latte* eliminates the need for application-specific re-implementation of common metrics, allowing consistent and reproducible model evaluation regardless of the deep learning framework of choice.

#### CRediT authorship contribution statement

**Karn N. Watcharasupat:** Conceptualization, Software, Validation, Resources, Data curation, Writing – original draft, Writing – review & editing, Project administration. **Junyoung Lee:** Software, Validation, Data curation. **Alexander Lerch:** Conceptualization, Validation, Resources, Writing – review & editing, Supervision.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgments

J. Lee acknowledges the support from the CN Yang Scholars Programme, Nanyang Technological University, Singapore. Part of this work was done while K. N. Watcharasupat was also supported by the CN Yang Scholars Programme, Nanyang Technological University, Singapore.

#### References

- [1] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, A. Lerchner,  $\beta$ -VAE: Learning basic visual concepts with a constrained variational framework, in: Proceedings Of The 5th International Conference On Learning Representations, 2017.
- [2] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, P. Abbeel, InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets, in: Proceedings Of The 30th Conference On Neural Information Processing Systems, 2016, pp. 2180–2188.
- [3] A. Pati, A. Lerch, Attribute-based regularization of latent spaces for variational auto-encoders, *Neural Comput. Appl.* 33 (9) (2020) 4429–4444.
- [4] H.H. Tan, D. Herremans, Music FaderNets: Controllable music generation based on high-level features via low-level feature modelling, in: Proceedings Of The 21st International Society For Music Information Retrieval Conference, 2020.
- [5] A. Pati, A. Lerch, Is disentanglement enough? On latent representations for controllable music generation, in: Proceedings Of The 22nd International Society For Music Information Retrieval Conference, 2021.
- [6] E. Vincent, R. Gribonval, C. Févotte, Performance measurement in blind audio source separation, *IEEE Trans. Audio Speech Lang. Process.* 14 (4) (2006) 1462–1469.
- [7] F.R. Stöter, A. Liutkus, D. Samuel, L. Miner, F. Voituret, Museval 0.4.0, 2021, GitHub, <https://github.com/sigsep/sigsep-mus-eval>. (Accessed 19 December 2021).
- [8] F. Locatello, S. Bauer, M. Lucic, G. Rätsch, S. Gelly, B. Schölkopf, O. Bachem, Challenging common assumptions in the unsupervised learning of disentangled representations, in: Proceedings Of The 36th International Conference On Machine Learning, 2019.
- [9] S. van Steenkiste, F. Locatello, J. Schmidhuber, O. Bachem, Are disentangled representations helpful for abstract visual reasoning? in: Proceedings Of The 33rd Conference On Neural Information Processing Systems, 2019.
- [10] F. Locatello, G. Abbati, T. Rainforth, S. Bauer, B. Schölkopf, O. Bachem, On the fairness of disentangled representations, in: Proceedings Of The 33rd Conference On Neural Information Processing Systems, 2019.
- [11] S. Duan, L. Matthey, A. Saraiva, N. Watters, C.P. Burgess, A. Lerchner, I. Higgins, Unsupervised model selection for variational disentangled representation learning, in: Proceedings Of The 8th International Conference On Learning Representations, 2019.
- [12] R. Shu, Y. Chen, A. Kumar, S. Ermon, B. Poole, Weakly supervised disentanglement with guarantees, in: Proceedings Of The 8th International Conference On Learning Representations, 2020.
- [13] F. Locatello, M. Tschannen, S. Bauer, G. Rätsch, B. Schölkopf, O. Bachem, Disentangling factors of variation using few labels, in: Proceedings Of The 8th International Conference On Learning Representations, 2020.
- [14] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D.G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, X. Zheng, TensorFlow: A System for large-scale machine learning, in: Proceedings Of The 12th USENIX Symposium On Operating Systems Design And Implementation, 2016, pp. 265–283.
- [15] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, in: Proceedings Of The 33rd Conference On Neural Information Processing Systems, 2019.
- [16] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N.J. Smith, R. Kern, M. Picus, S. Hoyer, M.H. van Kerkwijk, M. Brett, A. Haldane, J.F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T.E. Oliphant, Array programming with numPy, *Nature* 585 (7825) (2020) 357–362.
- [17] 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, E. Duchesnay, Scikit-learn: Machine learning in python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [18] The PyTorchLightning Team, TorchMetrics: Machine learning metrics for distributed, scalable PyTorch applications, 2020, GitHub, <https://github.com/PyTorchLightning/metrics>. (Accessed 19 December 2021).
- [19] W. Falcon, The PyTorch Lightning Team, Pytorch Lightning, 2019, GitHub, <https://github.com/PyTorchLightning/pytorch-lightning>. (Accessed 19 December 2021).

- [20] A. Kumar, P. Sattigeri, A. Balakrishnan, Variational inference of disentangled latent concepts from unlabeled observations, in: Proceedings Of The 6th International Conference On Learning Representations, 2018.
- [21] T.Q. Chen, X. Li, R. Grosse, D. Duvenaud, Isolating sources of disentanglement in variational autoencoders, in: Proceedings Of The 32nd International Conference On Neural Information Processing Systems, 2018.
- [22] K. Ridgeway, M.C. Mozer, Learning deep disentangled embeddings with the F-statistic loss, in: Proceedings Of The 32nd International Conference On Neural Information Processing Systems, 2018, pp. 185–194.
- [23] K.N. Watcharasupat, A. Lerch, Evaluation of latent space disentanglement in the presence of interdependent attributes, in: Extended Abstracts Of The Late-Breaking Demo Session Of The 22nd International Society For Music Information Retrieval Conference, 2021.
- [24] K.N. Watcharasupat, Controllable music: supervised learning of disentangled representations for music generation, 2021, Nanyang Technological University, Singapore.
- [25] H. Kim, A. Mnih, Disentangling by factorising, in: Proceedings Of The 35th International Conference On Machine Learning, 2018, pp. 4153–4171.
- [26] C. Eastwood, C.K. Williams, A framework for the quantitative evaluation of disentangled representations, in: Proceedings Of The 6th International Conference On Learning Representations, 2018.
- [27] R. Suter, D. Miladinović, B. Schölkopf, S. Bauer, Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness, in: Proceedings Of The 36th International Conference On Machine Learning, 2019, pp. 10593–10602.
- [28] M.F. Naeem, S.J. Oh, Y. Uh, Y. Choi, J. Yoo, Reliable fidelity and diversity metrics for generative models, in: Proceedings Of The 37th International Conference On Machine Learning, 2020, pp. 7133–7142.
- [29] V. Fomin, J. Anmol, S. Desroziers, J. Kriss, A. Tejani, High-level library to help with training neural networks in PyTorch, 2020, GitHub, <https://github.com/pytorch/ignite>. (Accessed 19 December 2021).
- [30] A. Hayakawa, M. Ishii, Y. Kobayashi, A. Nakamura, T. Narihira, Y. Obuchi, A. Shin, T. Yashima, K. Yoshiyama, Neural Network Libraries: A deep learning framework designed from engineers' perspectives, 2021, arXiv.