

Resilient VAE: Unsupervised Anomaly Detection at the SLAC Linac Coherent Light Source

Ryan Humble^{1,*}, William Colucho², Finn O’Shea², Daniel Ratner², and Eric Darve^{1,3}

¹Institute for Computational and Mathematical Engineering, Stanford University

²SLAC National Laboratory

³Department of Mechanical Engineering, Stanford University

Abstract. Significant advances in utilizing deep learning for anomaly detection have been made in recent years. However, these methods largely assume the existence of a normal training set (i.e., uncontaminated by anomalies) or even a completely labeled training set. In many complex engineering systems, such as particle accelerators, labels are sparse and expensive; in order to perform anomaly detection in these cases, we must drop these assumptions and utilize a completely unsupervised method. This paper introduces the Resilient Variational Autoencoder (ResVAE), a deep generative model specifically designed for anomaly detection. ResVAE exhibits resilience to anomalies present in the training data and provides feature-level anomaly attribution. During the training process, ResVAE learns the anomaly probability for each sample as well as each individual feature, utilizing these probabilities to effectively disregard anomalous examples in the training data. We apply our proposed method to detect anomalies in the accelerator status at the SLAC Linac Coherent Light Source (LCLS). By utilizing shot-to-shot data from the beam position monitoring system, we demonstrate the exceptional capability of ResVAE in identifying various types of anomalies that are visible in the accelerator.

1 Introduction

Anomaly detection (AD), which is the task of finding abnormal data or events, is a critical task for nearly all data-heavy, complex systems, such as industrial facilities, manufacturing, and large-scale science experiments [1–4]. These systems can produce thousands of real-time signals, which quickly overwhelm human operators that seek to monitor system performance. Identifying failures in these systems is a critical task, as failures can result in faulty outputs or cause damage to components. Unfortunately, the same complexity that overwhelms operators ensures that labeled data is rare or nonexistent and expensive to acquire.

In this work, we focus on detecting anomalies at SLAC’s Linac Coherent Light Source (LCLS), which is a free-electron laser (FEL) system that enables users to take X-ray snapshots of microscopic phenomena. Since LCLS loses approximately 3% of availability annually to unplanned downtime and experiences additional beam degradation without downtime, the task of identifying failures is critical to maximizing the amount of stable X-ray beam delivered to user experiments and thereby maximizing the scientific contribution of LCLS. Unfortunately,

*e-mail: ryhumble@stanford.edu

detecting anomalies at LCLS is challenging because of the data volume and lack of labeled data. This combination makes most traditional AD methods [5–8] challenging to apply. Deep learning approaches [9–15] are often quite effective in this high-volume, high-dimensional setting; however, they nearly always assume a normal-only training set. Since we lack labels, our training data will inevitably be contaminated by anomalies.

Our work’s main contributions are: (i) introduction of a new deep generative model for anomaly detection, called Resilient VAE, to cope with training data contamination (Section 2); and (ii) its application to identifying anomalies in accelerator status at LCLS (Section 3).

2 Resilient VAE

We consider a particular type of anomaly detection task where the data is high-dimensional and completely unlabeled. We suppose a dataset $\mathcal{D} = \{x\}$, where each sample is $x \in \mathbb{R}^D$ and D is the number of features. We suppose an (unknown) anomaly process that contaminates samples and features there-within, in our dataset. In particular, we assume that any training dataset has already been contaminated by anomalies, so there is no normal-only training set.

2.1 Background: Variational autoencoders for anomaly detection

Variational autoencoders (VAEs) [16] are probabilistic generative models designed for efficient inference and learning. The data is modeled as $p(x) = \mathbb{E}_{p(z)}[p_\theta(x|z)]$ where $p(z)$ is a prior over the latent representation $z \in \mathbb{R}^K$ and $p_\theta(x|z)$ is the *decoder* model. Typically $K < D$, so the latent space acts as an information bottleneck, thereby implicitly assuming our data has a low-dimension representation. When $p_\theta(x|z)$ is parameterized as a deep neural network, both the marginal likelihood $p(x)$ and posterior likelihood $p(z|x)$ are typically intractable [16]; thus, approximate variational inference is required, via an *encoder* model $q_\phi(z|x)$. The resulting variational bound on the marginal log-likelihood, called the evidence lower bound (ELBO), is

$$\log p(x) \geq \mathcal{L}_{\text{VAE}}(x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x)||p(z)), \quad (1)$$

and the lower bound $\mathcal{L}_{\text{VAE}}(x)$ is maximized w.r.t. to the parameters θ, ϕ . We refer to the first term, $\mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)]$, as the reconstruction likelihood, and the second term, $D_{KL}(q_\phi(z|x)||p(z))$, as the prior-matching penalty.

To apply the VAE model to anomaly detection, prior work most commonly uses the reconstruction likelihood to identify anomalies, where anomalies are assumed to be poorly reconstructed [17]. Furthermore, previous studies usually train the VAE model exclusively on normal data, as otherwise, the VAE might learn to reconstruct anomalous examples present in the training data. In fact, the VAE loss is weighted towards the anomalous examples, as both the reconstruction likelihood and prior-matching penalties are analogous to minimizing an L_2 norm, which is not robust to outliers.

2.2 Mixture model

Since the standard VAE is not robust to anomalies and we assume our training data is contaminated by anomalies, we propose an extension of the VAE model, which we call Resilient VAE (ResVAE), that aims to distinguish between normal and anomalous samples during training using a mixture model approach.

We begin by adding several binary latent variables to the standard VAE: a sample mixing variable $v \in \{0, 1\}$ and feature mixing variables $w_d \in \{0, 1\}$, where 1 means inlier and 0 outlier.

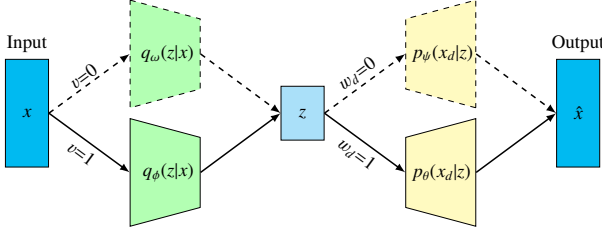


Figure 1. Network diagram for our proposed ResVAE. ResVAE adds a sample mixing variable v and feature mixing variables w_d and splits the encoder and decoders into an inlier (solid path) and outlier (dotted path) version.

These mixing variables control whether an inlier or outlier path through the ResVAE model is taken, as shown in Figure 1. With the new latent variables, our data is now modeled as $p(x) = \mathbb{E}_{p(v,w,z)}[p(x|v,w,z)]$; we use the variational posterior $q(v,w,z|x)$ for inference. Our model supposes Bernoulli distributions for the mixing variables, a mixture prior on the z latent variable, and encoder and decoder mixture models:

$$p(v) = \text{Ber}(v; \rho), \quad q_\gamma(v|x) = \text{Ber}(v; \gamma(x)), \quad (2)$$

$$p(w_d|v) = \text{Ber}(w_d; \alpha_d)^v \text{Ber}(w_d; 0)^{1-v}, \quad q_\pi(w_d|v, x) = \text{Ber}(w_d; \pi_d(x))^v \text{Ber}(w_d; 0)^{1-v}, \quad (3)$$

$$p(z|v) = p(z)^v p_o(z)^{1-v}, \quad q(z|v, x) = q_\phi(z|x)^v q_\omega(z|x)^{1-v}, \quad (4)$$

$$p(x|v, w, z) = \prod_{d=1}^D p_\theta(x_d|z)^{vw_d} p_\psi(x_d|z)^{1-vw_d}, \quad (5)$$

where $\rho, \alpha_d \in [0, 1]$ can be used to capture prior knowledge on the degree of contamination in the training data. We revisit these prior parameters in Section 2.4. The functions $\gamma(x)$ and $\pi_d(x)$ yield estimates of the mixing variables v and w_d . As we show in Section 2.3, they have simple closed-form expressions as functions of the encoder and decoder parameters; this avoids the need for γ and π_d to be parameterized as separate neural networks.

Our mixture model introduced several outlier distributions— $p_o(z)$, $p_\psi(x|z)$, and $q_\omega(z|x)$ —whose purpose is to provide an alternate model for outlier data, thereby allowing the inlier encoder $q_\phi(z|x)$ and decoder $p_\theta(x|z)$ to focus on representing only the inlier data in our training data. Concretely, we define the inlier and outlier distributions as

$$\begin{aligned} p(z) &= \mathcal{N}(z; 0, I), & p_\theta(x_d|z) &= p(x_d; \eta_d(z; \theta)), & q_\phi(z|x) &= \mathcal{N}(z; \mu(x; \phi), \Sigma(x; \phi)), \\ p_o(z) &= \mathcal{N}(z; 0, \delta_z^2 I), & p_\psi(x_d|z) &= p(x_d; \eta_d(z; \theta_0) / \delta_x^2), & q_\omega(z|x) &= q_{\phi_0}(z|x), \end{aligned}$$

where δ_x, δ_z are hyperparameters of the method, $p_\theta(x_d|z)$ is taken to be an exponential family distribution with natural parameters $\eta_d(z; \theta)$, θ_0 is a copy of θ , and ϕ_0 is a copy of ϕ (where the copy prevents gradients via automatic differentiation). (We refer the reader to a longer version of the paper [18] for an explanation of these modeling choices and subsequent derivations.)

2.3 ResVAE loss and coordinate ascent

Instead of maximizing the ELBO, we adopt the constrained-optimization approach of β -VAE [19], which introduces several coefficients regularizing the prior-matching penalties. Our objective is then to maximize the following loss function:

$$\begin{aligned} \max_{\phi, \theta, \omega, \psi} \mathcal{L}(x) &= \mathbb{E}_{q(v,w|x)} \left[\mathbb{E}_{q(z|v,x)} [\log p(x|v,w,z)] \right] - \beta_1 \mathbb{E}_{q_\gamma(v|x)} [D_{KL}(q(z|x, v) \| p(z|v))] \\ &\quad - \sum_{d=1}^D \frac{1}{\beta_{2,d}} \mathbb{E}_{q_\pi(w_d|v,x)} [D_{KL}(q_{\pi_d}(w_d|v, x) \| p(w_d|v))] - \frac{1}{\beta_3} D_{KL}(q_\gamma(v|x) \| p(v)). \end{aligned}$$

β_1 plays an identical role to that in β -VAE [19], controlling the latent information bottleneck against the reconstruction objective. We discuss the other regularization coefficients $\beta_{2,d}, \beta_3$ in Section 2.4.

We can simplify this loss function by performing coordinate ascent (as in [20, 21]) on $\gamma(x)$ and $\pi_d(x)$ since $\mathcal{L}(x)$ is convex in these variables for each x . Their optimal values are:

$$\pi_d(x) = \sigma(\beta_{2,d}r_d(x) + \text{logit}(\alpha_d)), \quad \gamma(x) = \sigma(\beta_3g(x) + \text{logit}(\rho)), \quad (6)$$

where $\sigma(\cdot)$ is the sigmoid function, $\text{logit}(\cdot)$ is the logit function, and $r_d(x)$ and $g(x)$ can be interpreted as likelihood ratio tests [22] (i.e., $r_d(x) > 0$ suggests x_d is drawn from the inlier model instead of the outlier model and the opposite for $r_d(x) < 0$).

Using the closed form for $\pi_d(x)$ and $\gamma(x)$ and the given outlier distributions, and removing untrainable terms, our ResVAE loss function becomes:

$$\max_{\phi, \theta} \mathcal{L}_{\text{ResVAE}}(x) = \gamma(x) \left[\sum_d \pi_d(x) \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x_d|z)] - \beta_1 D_{\text{KL}}(q_\phi(z|x) \| p(z)) \right]. \quad (7)$$

This is nearly the form of the standard VAE loss (or β -VAE loss [19]), with the addition of several weighting factors that either preserve the sample (or features) in the loss function or remove it. $\gamma(x)$ weights entire samples and $\pi_d(x)$ weights individually features.

2.4 Learnable logistic parameters

We can interpret the form of $\pi_d(x)$ and $\gamma(x)$ as a logistic regression, where the input variables are the likelihood ratios $r_d(x)$ and $g(x)$ respectively and the parameters are the regularization coefficients $\beta_{2,d}, \beta_3$ and contamination prior parameters α_d, ρ . Ideally, $\pi_d(x)$ and $\gamma(x)$ should be nearly 1 for inliers and 0 for outliers. Achieving this requires setting the logistic parameters appropriately.

To ensure at all stages of training that $\pi_d(x)$ and $\gamma(x)$ are expressive, we propose to learn the logistic parameters. Since we can interpret $\pi_d(x)$ and $\gamma(x)$ as discriminators (as in generative adversarial nets [23]), we train them as binary classifiers with the binary cross entropy loss:

$$\max_{\beta_{2,d}, \alpha_d} \mathcal{L}_{\pi_d}(x) = \mathbb{E}_{p(w_d|v=1, x)} [\log q_\pi(w_d|v=1, x)], \quad \max_{\beta_3, \rho} \mathcal{L}_\gamma(x) = \mathbb{E}_{p(v|x)} [\log q_\gamma(v|x)],$$

where $q_\pi(w_d|v=1, x)$ depends on the logistic parameters $\beta_{2,d}, \alpha_d$ through the definition of $\pi_d(x)$ and $q_\gamma(v|x)$ similarly depends on β_3, ρ through $\gamma(x)$. The label distributions of the binary cross entropy terms are the posterior distributions $p(w_d|v=1, x) = \text{Ber}(w_d; \tilde{w}_d(x))$ and $p(v|x) = \text{Ber}(v; \tilde{v}(x))$. However, since these posteriors are intractable (which originally motivated the use of approximate posteriors), the labels $\tilde{w}_d(x)$ and $\tilde{v}(x)$ must be approximated.

In defining our label approximations, we want to achieve two goals. First, we want the empirical distributions of $\pi_d(x)$ and $\gamma(x)$ to match our prior knowledge of the dataset's contamination. Since ρ, α_d are now learned as logistic parameters instead of specified as contamination priors, we re-introduce the idea of a contamination prior through new prior label distributions $f_{w_d}(w_d)$ and $f_v(v)$ over the entire dataset. Second, since we assume that inlier examples have higher likelihood ratios $r_d(x)$ and $g(x)$ than outlier examples, we want $\pi_d(x), \gamma(x)$ to be increasing in the likelihood ratios $r_d(x), g(x)$ respectively. Therefore, we approximate the labels as:

$$\tilde{w}_d(x) \approx F_{w_d} \left(\mathbb{P}_{y \sim \mathcal{D}} (\pi_d(x) \geq \pi_d(y)) \right), \quad \tilde{v}(x) \approx F_v \left(\mathbb{P}_{y \sim \mathcal{D}} (\gamma(x) \geq \gamma(y)) \right), \quad (8)$$

where the inner rank statistic terms ensure the property of increasing in likelihood ratio and the cumulative distribution functions (CDFs) $F_{w_d}(w_d)$ and $F_v(v)$ (of $f_{w_d}(w_d)$ and $f_v(v)$ respectively) ensure the prior-empirical label distribution matching.

3 Application to LCLS

We now apply our ResVAE method to detecting anomalies at LCLS. We consider pulse-by-pulse measurements of the beam, as measured by stripline beam position monitors (BPMs) [24] and recorded by the beam-synchronous acquisition system [25]. Each of the 150 BPMs records three signals (X, Y, and TMIT) for each electron pulse, where X and Y measure transverse position and TMIT (transmitted intensity) measures the passing beam charge, for a total of 450 signals. Since each pulse should be i.i.d. under normal operation, we currently consider each pulse to be a separate input; using windows of consecutive pulses is left for future work.

Since the operators of LCLS are constantly changing the machine settings, we limit our study to periods of “steady” operation at 120 Hz (as detailed in [26]) and train on a single day of data at a time. Nonetheless, due to the contamination in the training data, the standard VAE [16], RVAE [21], and DSVDD [15] methods, as well as ResVAE without logistic learning, fail to train; they experience disrupting gradients due to the anomalous examples. Gradient clipping does not fix the training failures, as this solely moderates the magnitudes of applied gradients instead of ignoring the gradient contribution of anomalous examples entirely. We train our ResVAE model using the PyTorch framework [27] and the Adam optimizer [28] using mini-batches. (Full details on the training setting and architecture are found in the longer version of our paper [18].)

To explore the anomalies detected by our method, we extract all anomalous pulses (where $\gamma(x) \leq 0.5$) and then perform clustering on the vector $\pi(x)$ (which provides feature-level attribution) using UMAP [29]. We show the existence of several anomaly clusters in Figure 2. We also dive into several clusters to determine the nature of the anomaly. Figure 3 shows an anomaly where the X-ray beam is partially lost, observed in the beam’s TMIT signals. Figure 4 displays the onset of an anomaly observed in the beam’s X signals. ResVAE is capable of providing operators with information regarding both the timing and location of the anomaly.

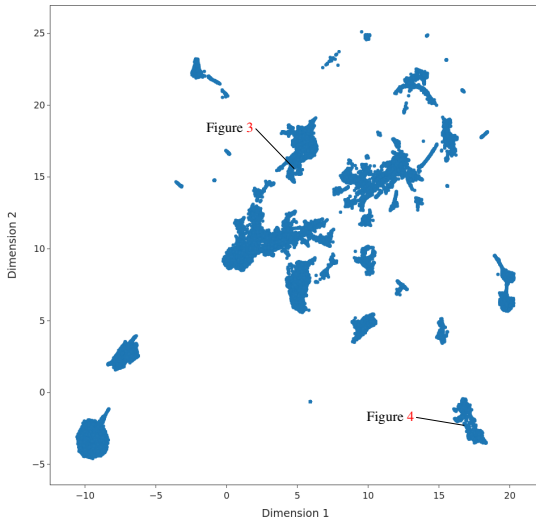


Figure 2. UMAP [29] clustering of the feature-level anomaly attribution vector $\pi(x)$ for anomalous pulses (where $\gamma(x) \leq 0.5$).

4 Related Work

There are numerous prior works studying the problem of unsupervised anomaly detection. Classical methods [5–8], typically rely on density or distance and therefore struggle with high-dimensional inputs, especially when anomalies are present in the training set. The main alternative methods rely on compression techniques, such as autoencoders [9, 10], VAEs [11],

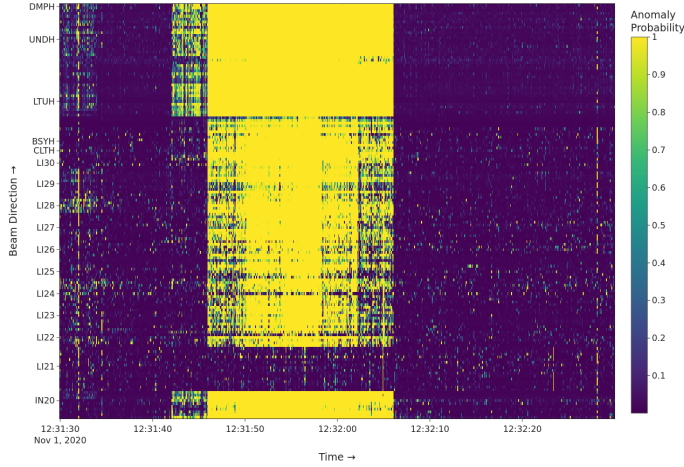


Figure 3. Example of a beam-loss anomaly detected by ResVAE and its feature-level anomaly attribution (where the feature-level anomaly probability is $1 - \pi_d(x)$). We only show the TMIT signals.

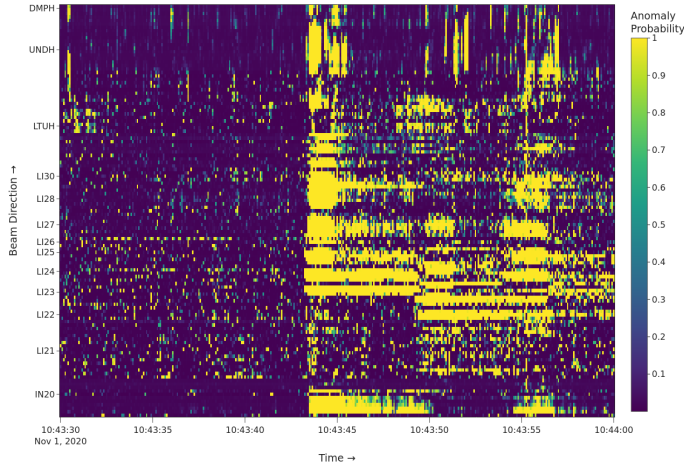


Figure 4. Example of anomaly detected by ResVAE and its feature-level anomaly attribution (where the feature-level anomaly probability is $1 - \pi_d(x)$). We only show the X signals.

and generative adversarial nets (GANs) [12, 13], or combine deep compression with classical methods [14, 15]. Nonetheless, these methods almost always assume a normal training set to avoid learning to compress anomalous examples.

Several methods attempt to add robustness to outliers to existing methods, including robust PCA [30, 31], β -robust VAE [32], and RVAE [21]. However, all of these approaches have various drawbacks. Although robust PCA (and its deep counterpart) can effectively identify both sample and feature anomalies in contaminated training data, its main drawbacks are the relative difficulty of interpreting the hyperparameter controlling the outlier fraction and the need to store the outlier matrix S of size equal to the dataset. β -robust VAE [32] replaces the reconstruction likelihood term of the traditional ELBO with a more robust β -divergence term; however, this does not address robustness concerns in the latent space nor does this enable the model to identify feature outliers. Our approach is most similar to and extends the RVAE [21] work, following its general structure and derivation but making several key changes to add robustness to the latent space and ensure expressive weights.

5 Conclusion

This study introduces the ResVAE anomaly detection method, an innovative extension of the VAE model [16] designed to detect anomalies even when trained on contaminated data. We

provide a theoretical explanation for the limitations of VAE when faced with contaminated training data and propose a solution using a mixture model that assigns weights to both samples and features during training, effectively mitigating the impact of outliers. Subsequently, we apply this method to anomaly detection at LCLS, successfully identifying anomalous accelerator status and illustrating various types of anomalies.

References

- [1] J.P. Edelen, N.M. Cook, arXiv preprint arXiv:2112.07793 (2021)
- [2] M. Radaideh, C. Pappas, P. Ramuhalli, S. Cousineau, Annual Conference of the PHM Society **14** (2022)
- [3] E. Fol, R. Tomás, J.C. de Portugal, G. Franchetti, Physical Review Accelerators and Beams **23** (2020)
- [4] G. Valentino, R. Bruce, S. Redaelli, R. Rossi, P. Theodoropoulos, S. Jaster-Merz, *Anomaly detection for beam loss maps in the Large Hadron Collider*, in *Journal of Physics: Conference Series* (IOP Publishing, 2017), Vol. 874, p. 012002
- [5] D.A. Reynolds et al., Encyclopedia of biometrics **741** (2009)
- [6] E. Parzen, The Annals of Mathematical Statistics **33**, 1065 (1962)
- [7] B. Schölkopf, J.C. Platt, J. Shawe-Taylor, A.J. Smola, R.C. Williamson, Neural Computation **13**, 1443 (2001)
- [8] F.T. Liu, K.M. Ting, Z.H. Zhou, *Isolation Forest*, in *2008 Eighth IEEE International Conference on Data Mining* (IEEE, 2008)
- [9] J. Chen, S. Sathe, C. Aggarwal, D. Turaga, *Outlier detection with autoencoder ensembles*, in *Proceedings of the 2017 SIAM international conference on data mining* (SIAM, 2017)
- [10] W. Lu, Y. Cheng, C. Xiao, S. Chang, S. Huang, B. Liang, T.S. Huang, IEEE Trans. Image Process. **26**, 4321 (2017)
- [11] Y. Wang, B. Dai, G. Hua, J.A.D. Aston, D.P. Wipf, *Green Generative Modeling: Recycling Dirty Data using Recurrent Variational Autoencoders*, in *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*, edited by G. Elidan, K. Kersting, A. Ihler (AUAI Press, 2017)
- [12] Y. Su, Y. Zhao, C. Niu, R. Liu, W. Sun, D. Pei, *Robust anomaly detection for multivariate time series through stochastic recurrent neural network*, in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining* (2019)
- [13] Y. Liu, Z. Li, C. Zhou, Y. Jiang, J. Sun, M. Wang, X. He, IEEE Transactions on Knowledge and Data Engineering **32**, 1517 (2019)
- [14] B. Zong, Q. Song, M.R. Min, W. Cheng, C. Lumezanu, D. Cho, H. Chen, *Deep Autoencoding Gaussian Mixture Model for Unsupervised Anomaly Detection*, in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings* (OpenReview.net, 2018)
- [15] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S.A. Siddiqui, A. Binder, E. Müller, M. Kloft, *Deep One-Class Classification*, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy, A. Krause (PMLR, 2018), Vol. 80 of *Proceedings of Machine Learning Research*, pp. 4393–4402
- [16] D.P. Kingma, M. Welling, *Auto-Encoding Variational Bayes*, in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, edited by Y. Bengio, Y. LeCun (2014)
- [17] J. An, S. Cho, Special lecture on IE **2**, 1 (2015)

- [18] R. Humble, W. Colocho, F. O'Shea, D. Ratner, E. Darve, *Resilient vae: Unsupervised anomaly detection at the slac linac coherent light source* (2023), 2309.02333
- [19] I. Higgins, L. Matthey, A. Pal, C.P. Burgess, X. Glorot, M.M. Botvinick, S. Mohamed, A. Lerchner, *beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework*, in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings* (OpenReview.net, 2017)
- [20] M.I. Jordan, Z. Ghahramani, T.S. Jaakkola, L.K. Saul, *Machine learning* **37**, 183 (1999)
- [21] S. Eduardo, A. Nazabal, C.K.I. Williams, C. Sutton, *Robust Variational Autoencoders for Outlier Detection and Repair of Mixed-Type Data*, in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, edited by S. Chiappa, R. Calandra (PMLR, 2020), Vol. 108 of *Proceedings of Machine Learning Research*
- [22] H.L. Van Trees, *Detection, estimation, and modulation theory, part I: detection, estimation, and linear modulation theory* (John Wiley & Sons, 2004)
- [23] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A.C. Courville, Y. Bengio, *Generative Adversarial Nets*, in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, edited by Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, K.Q. Weinberger (2014), pp. 2672–2680
- [24] T. Straumann, R. Akre, R.G. Johnson, D. Kotturi, P. Krejcik, E. Medvedko, J. Olsen, S. Smith, *LCLS Beam-Position Monitor Data Acquisition System*, in *2007 International Conference on Accelerator and Large Experimental Physics Control Systems (JACoW, Knoxville, Tennessee, 2007)*
- [25] K. Kim, S. Allison, T. Straumann, E. Williams, *Real-Time Performance Improvements and Consideration of Parallel Processing for Beam Synchronous Acquisition (BSA)*, in *Proc. of International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'15) Melbourne Australia 17-23 October 2015 (JACoW, Geneva, Switzerland, 2015)*, Number 15 in *International Conference on Accelerator and Large Experimental Physics Control Systems*, pp. 992–994, ISBN 978-3-95450-148-9, doi:10.18429/JACoW-ICALEPCS2015-WEPGF122
- [26] R. Humble, F.H. O'Shea, W. Colocho, M. Gibbs, H. Chaffee, E. Darve, D. Ratner, *Phys. Rev. Accel. Beams* **25**, 122804 (2022)
- [27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Curran Associates, Inc., 2019), Vol. 32
- [28] D.P. Kingma, J. Ba, *Adam: A Method for Stochastic Optimization*, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, edited by Y. Bengio, Y. LeCun (2015)
- [29] L. McInnes, J. Healy, N. Saul, L. Großberger, *Journal of Open Source Software* **3**, 861 (2018)
- [30] E.J. Candès, X. Li, Y. Ma, J. Wright, *J. ACM* **58**, 11:1 (2011)
- [31] C. Zhou, R.C. Paffenroth, *Anomaly Detection with Robust Deep Autoencoders*, in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017 (ACM, 2017)*, pp. 665–674
- [32] H. Akrami, A.A. Joshi, J. Li, S. Aydıre, R.M. Leahy, *Knowl. Based Syst.* **238**, 107886 (2022)