

---

# Normalizing flows for novelty detection in industrial time series data

---

Maximilian Schmidt<sup>1</sup> Marko Simic<sup>1</sup>

## Abstract

Flow-based deep generative models learn data distributions by transforming a simple base distribution into a complex distribution via a set of invertible transformations. Due to the invertibility, such models can score unseen data samples by computing their exact likelihood under the learned distribution. This makes flow-based models a perfect tool for novelty detection, an anomaly detection technique where unseen data samples are classified as normal or abnormal by scoring them against a learned model of normal data. We show that normalizing flows can be used as novelty detectors in time series. Two flow-based models, Masked Autoregressive Flows and Free-form Jacobian of Reversible Dynamics restricted by autoregressive MADE networks, are tested on synthetic data and motor current data from an industrial machine and achieve good results, outperforming a conventional novelty detection method, the Local Outlier Factor.

## 1. Introduction

Novelty detection comprises anomaly detection algorithms which first learn a model of the normal data, and at inference time compute a novelty score of unseen data samples under the learned model. Data samples are flagged as normal or abnormal by comparing their novelty score to a learned decision boundary. In that sense, novelty detection can be seen as a one-class classification task. One important characteristic, and appealing property, of novelty detection algorithms is that during training time they do not require access to instances of anomalous data. This is important in situations in which we have little or no anomalous data, and it is hard and costly to obtain such data. Because of this, novelty detection is widely used in many domains, such as medical diagnostics (Tarassenko et al., 1995), security of electronic systems (Patcha & Park,

2007), or mobile robotics (Hornung et al., 2014).

Pimentel et al. (2014) classify novelty detection algorithms in several groups: probabilistic, distance-based, reconstruction-based, domain-based, and information-theoretic based techniques. For our application, we initially chose four methods that belong to different algorithm families outlined above: support vector machine (SVM) (Schölkopf et al., 2000), isolation forest (IF) (Liu et al., 2012), local outlier factor detector (LOF) (Breunig et al., 2000) as representative of conventional machine learning algorithms, and normalizing flows (NF) (Kingma et al., 2016) as representative of deep learning algorithms. After preliminary tests, out of three mentioned conventional methods, LOF yielded best classification accuracy, so we focused our further effort on LOF and normalizing flows.

Normalizing flows are a class of deep generative models that leverage invertible neural networks to learn a mapping between a simple base distribution and a given data distribution. The invertibility allows for two important use cases: generation of new data and classification of input data samples by computing the likelihood of such samples. The latter one makes them a suitable candidate algorithm for novelty detection. Compared to classical algorithms, normalizing flows allow us to flexibly constrain the learned distributions, for instance by enforcing autoregressive property when modeling time series.

In this work, we demonstrate the applicability of normalizing flows for novelty detection in time series. We apply two different flow models, masked autoregressive flows (MAF) (Papamakarios et al., 2017) and FFJORD (Grathwohl et al., 2019) restricted by a Masked Autoencoder for Distribution Estimation (MADE) architecture (Germain et al., 2015) to synthetic data and motor current time series data from an industrial machine. Both flow-based models achieve superior results over the local outlier factor method. Furthermore, we demonstrate the generation of new data samples with the learned flow models which can give rise to further use cases in the domain of anomaly detection and defect analysis in industrial machines.

---

<sup>1</sup>Ascent Robotics Inc., Japan. Correspondence to: Maximilian Schmidt <maximilian@ascent.ai>.

## 2. Datasets

We create synthetic data to test the ability of the models for novelty detection by generating random time series with a defined autocorrelation function. We specify the autocorrelation function as  $f(\Delta t) = \exp(-|\Delta t|/\tau) \cos(\Delta t/15)$  which defines the covariance matrix  $\Sigma$  for a given length of time series  $T$  and compute its Cholesky decomposition  $B$  such that  $\Sigma = BB^T$ . We then generate data samples by drawing white noise samples  $y \sim N(0, 1) \in \mathbb{R}^T$  and transforming them to samples of our time series with given autocorrelation as  $z = yB$ . We define normal samples to have  $\tau = 50$  and vary the decay time to create abnormal samples. By construction, the inter-sample mean at each timestep is 0. To keep the variance across samples at every time step equal between normal and abnormal samples, we divide abnormal samples  $y'$  by their inter-sample variance and multiply by the variance of the normal time series:  $y'_i(t) \leftarrow y'_i(t)/\text{Var}_j(y'_j(t)) \cdot \text{Var}_j(y_j(t))$ .

To test on real data, we used a dataset that was generously provided by Kawasaki Heavy Industries Ltd (KHI). Data is an electric motor current signal (measured in Amperes), collected from the electrical motor of one of KHI’s products. Dataset contains both the signal during motor’s normal operation, and anomalous signal after a gearbox connected to the motor experienced an undefined problem in its operation. There are eight different patterns of motor operation, each pattern lasting between 5 and 30 seconds. Signals were sampled with frequency of 500 Hz. Ratio of normal to anomalous data in the dataset was roughly equal, but for model training purposes only a subset of normal data was used, while held-out normal data and anomalous data were used for testing purposes. One sample of normal and anomalous data is shown in Figure 3A. In this work, we use a subset of one of eight patterns that, based on classification performance of the LOF, we found to be the most challenging pattern for anomaly detection.

## 3. Model setup

We train three different models on our data: Masked Autoregressive Flows, Continuous Normalizing Flows using Free-form Jacobian of Reversible Dynamics with a MADE network, and Local Outlier Factor.

- MAF uses a stack of fixed number of affine layers whose scale and shift parameters are computed by an autoregressive network, here implemented using the MADE architecture (Germain et al., 2015). Given a latent random variable  $z_t \sim \pi(z_t)$ , the transformed variable  $x_t$  is computed as  $x_t = z_t \odot \mu(\mathbf{x}_{1:t-1}) + \sigma(\mathbf{x}_{1:t-1}) \sim p(x_t|\mathbf{x}_{1:t-1})$ , where the scale and shift terms  $\mu, \sigma$  are efficiently computed by one forward pass through a MADE network. We chose MAF over

Inverse Autoregressive Flows (IAF) (Kingma et al., 2016) because MAF offers fast evaluation of data likelihood which is essential in novelty detection. IAF, on the other hand, offers fast generation of new data but slow evaluation of test data. In our experiments we chose the standard normal distribution as the base distribution  $z \sim N(0, I)$  and stack 5 coupling layers each with MADE networks consisting of 3 hidden layers with 256 units each and tanh activation function.

- The FFIJORD model (Grathwohl et al., 2019) extends the idea of continuous normalizing flows (CNF) (Chen et al., 2018) by an improved estimator of the log-density of samples. CNF models the latent variable  $z$  with an ordinary differential equation  $\partial z / \partial \tilde{t} = f(z, \tilde{t})$  so that transforming from latent to data space is equivalent to integrating the ODE from pseudo times  $\tilde{t}_0$  to  $\tilde{t}_1$ :  $x = z(t_1) = \int_{\tilde{t}_0}^{\tilde{t}_1} f(z, t) dt$  (see Grathwohl et al. (2019) for details). The function  $f$  is represented by a neural network. We here chose a MADE architecture to enforce the autoregressive property between different time samples:  $f(z_t, \tilde{t}) = f(z_{1:t-1}, \tilde{t})$ . In our experiments, we use 2 hidden layers with 256 neurons and tanh activation function. More details on the experiments in Section 7.
- Local outlier factor is a distance based novelty detection algorithm that assigns a degree of being an outlier to each data point. This degree, local outlier factor, is determined by comparing the local density of a data point to the local density of its neighboring points. A point that has significantly lower local density than its neighbors is considered to be an outlier. The main parameter used to influence the algorithm’s performance is *MinPts* which specifies the number of points to be considered as neighborhood of a datapoint  $p$ . Breunig et al. (2000) establishes that inliers have value of LOF equal to 1, while LOF for outliers is greater than 1. It also provides the tightness of the lower and upper bound for outlier’s LOF values.

## 4. Experiments

We test our idea of using flow models as novelty detectors for time series on a set of synthetic data where we control the deviation between normal and abnormal samples. The data is created as time series with a defined autocorrelation where we systematically vary the time constant of the autocorrelation (see Section 2 for details). Normal and abnormal data only vary in their correlation between time points, while the ensemble statistics (in terms of mean and variance across samples) are equal at every time step (Figure 1A, B). This makes anomaly detection challenging because the model cannot resort to simply representing mean and variance of time steps independently but rather needs

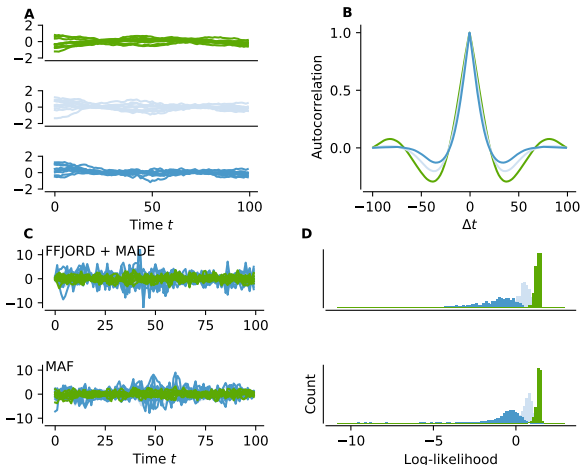


Figure 1. Flow models applied to synthetic data: normal data with  $\tau = 50$  (green) and abnormal data ( $\tau = 30$ , light blue and  $\tau = 20$ , dark blue). A: 10 random samples of the time series. B: Mean autocorrelation across 1000 samples. C: 10 samples transformed by applying FJORD+MADE and MAF. D: Histogram over mean likelihood values per time points of normal and abnormal samples.

to learn a joint distribution representing temporal correlations.

We train the two flow models on the normal data and then apply them to unseen normal samples as well as abnormal data samples (Figure 1C). Both models transform the normal samples to approx. white noise while the transformed abnormal samples clearly deviate from white noise. Consistent with the visual inspection of transformed samples, the models assign lower likelihood to abnormal samples than normal samples, and this difference becomes more clear with increasing deviation of abnormal samples, i.e. decreasing time constant of the autocorrelation function (Figure 1, dark blue vs. light blue in data points).

To use the trained model as a novelty detector, we define a decision boundary, i.e. a likelihood value which separates abnormal from normal data samples. To judge the quality of a novelty detector, a typical metric is the receiver operating characteristic (ROC) curve. Given a model and abnormal data samples, we vary the decision boundary and measure the rate of false positive (normal data classified as abnormal) versus true positive (abnormal data classified as abnormal). Thus, the steeper the slope of the ROC curve, the better. For varying time constants in the autocorrelation of the synthetic data, we compute the ROC curve of the two flow-based models and compare them to the local outlier factor (LOF) (Figure 2). Both flow-based models quickly deviate from the chance-level at  $\tau = 50$  (in this case, the test data is drawn from the same distribution as normal data) for decreasing time constant. The LOF model




Figure 2. Novelty detection results for synthetic data. Receiver operating characteristic (ROC) curves for all three tested models on abnormal synthetic data with decreasing time constants from 50 (light blue) to 20 (dark blue). The dashed line indicates chance level.

does not impose autoregressive constraints on the time series and thus, it performs worse in modelling the temporal correlations between data points in the time series.

We test the feasibility of the flow-based models on real data (see Section 2 for details). Both normal and abnormal time series follow a very similar global curve (Figure 3A, top) and are visually hardly distinguishable. We thus split normal samples into train and test data and subtract the mean across training samples from normal and abnormal samples (Figure 3A, middle and bottom). Abnormal and normal samples display different temporal correlations (Figure 3B), and consequently after training on the normal data both models clearly separate normal test data and abnormal data (Figure 3C). We find that the FJORD+MADE diverges when applied to abnormal data samples after the training has been performed for sufficient number of epochs (depending on the hyperparameters), where very low training and test loss has been reached. The ordinary differential equation becomes unstable for samples deviating from the training distribution. For the purposes of visualization, we thus decide to stop the training after 140 epochs for Figure 3. While the normal samples are approximately mapped to uncorrelated white noise, the transformed abnormal samples clearly deviate from white noise. To score the samples, we again compute the likelihood per time point and bin them into histograms (Figure 3D). It is trivial to define a decision boundary which reaches 0% false positive rate and 100% true positive rate. The LOF model reaches an accuracy of around 100% true positive rate and 0.5% false positive rate on this task. The flow-based models marginally outperform the simpler method.

The flow-based models enable us to generate new, artificial data. We draw random samples of white noise (Figure 4A, top) and pass them through the models to transform the random samples into samples from the data distribution (Figure 4A, bottom). The generated samples fit the empirical autocorrelation very well (Figure 4B), demonstrating that we indeed can generate artificial samples of normal data

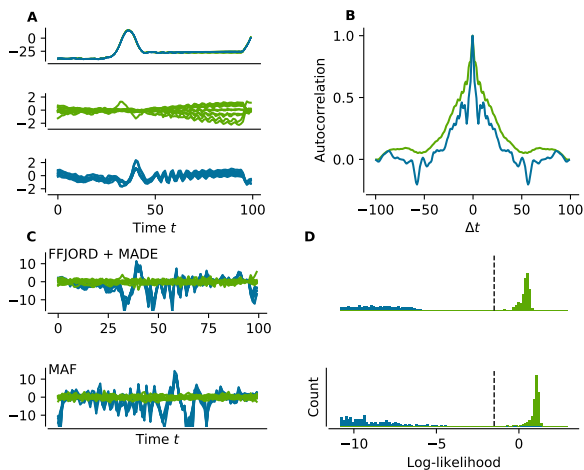


Figure 3. Flow models applied to real data: normal data (green) and abnormal data (dark blue). A: 10 samples of normal and abnormal data in the original domain (top) and centered with zero mean (middle and bottom). B: Mean autocorrelation across samples. C: Transformed samples by applying FFJORD+MADE and MAF. D: Histogram over likelihood values of normal and abnormal samples. The vertical dashed line indicates a possible decision boundary.

from the learned models.

## 5. Discussion

In this work, we evaluate the use of flow-based generative models for novelty detection in time series. Since normalizing flows approximate distribution of the data and score data samples based on their likelihood, they possess the sufficient ingredients for novelty detection. In comparison to other conventional methods, they can learn very flexible distributions and allow the modeler to impose helpful constraints such as autoregressive properties in time series.

To test our idea, we use two flow-based models: Masked Autoregressive Flows (MAF) and a Free-form Jacobian of Reversible Dynamics (FFJORD) continuous-time flow. We restrict the FFJORD model with autoregressive constraint imposed through the usage of autoregressive MADE networks, which to the best of our knowledge is the first time. We train the models on synthetic data, which we could control to make novelty detection challenging, and demonstrate good classification performance, outperforming a conventional method for novelty detection, Local Outlier Factor (LOF). Applied to less challenging real data from an industrial machine, the models reach perfect accuracy, marginally better than LOF.

As an extension of this work, the flow-based models could be trained to learn the transition from normal to abnormal samples. This could support the analysis of the defect caus-

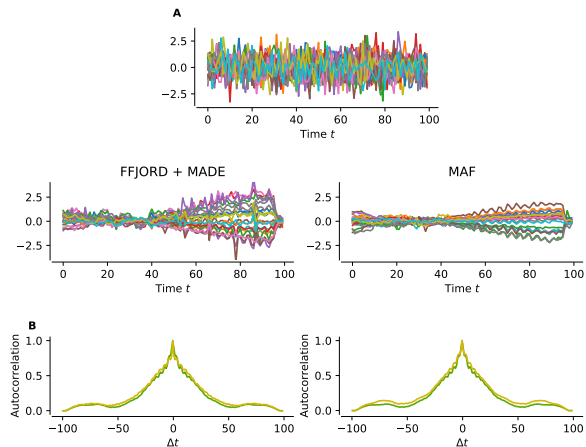


Figure 4. Generation of new data samples. (A) Samples of white noise ( $z(t) \sim N(0, 1) \forall t$ ) (top) are transformed into samples from the data distribution with the two trained models (bottom). (B) The autocorrelation averaged across 20 generated samples (yellow) matches the empirical autocorrelation of the data (green) very well. The FFJORD+MADE model was trained for 600 epochs.

ing the emergence of abnormal data in the industrial machine. Furthermore, the learned transition from normal to abnormal data could be applied to new motor operation patterns and generate abnormal data, thereby helping to understand the possible failure modes of the machine defect.

## 6. Acknowledgments

We would like to thank Jun Yamaguchi and Hitoshi Hasunuma from Kawasaki Heavy Industries Ltd. and Shunji Goto from Ascent Robotic Inc. for their support of this project.

## References

- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P., Horsfall, P., and Goodman, N. D. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine Learning Research*, 2018.
- Breunig, M. M., Kriegel, H. P., Ng, R. T., and Sander, J. LOF: Identifying density-based local outlier. In *ACM SIGMOD Int. Conf. on Management of Data*, pp. 93–104, 2000.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, pp. 6572–6583, 2018.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H.

- Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pp. 881–889, 2015.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. FFJORD: Free-Form continuous dynamics for scalable reversible generative models. pp. 1–13, 2019.
- Hornung, R., Urbanek, H., Klodmann, J., Osendorfer, C., and Van Der Smagt, P. Model-free robot anomaly detection. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3676–3683. IEEE, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pp. 4743–4751, 2016.
- Liu, F. T., Ting, K. M., and Zhou, Z. H. Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):3, 2012.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pp. 2338–2347, 2017.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- Patcha, A. and Park, J.-M. An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12):3448–3470, aug 2007. doi: 10.1016/j.comnet.2007.02.001.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Pimentel, M. A., Clifton, D. A., Clifton, L., and Tarassenko, L. A review of novelty detection. *Signal Processing*, 99:215–249, 2014.
- Schölkopf, B., Williamson, R. C., Smola, A. J., Shawe-Taylor, J., and Platt, J. C. Support vector method for novelty detection. In *Advances in neural information processing systems*, pp. 582–588, 2000.
- Tarassenko, L., Hayton, P., Cerneaz, N., and Brady, M. Novelty detection for the identification of masses in mammograms. In *Proceedings of the 4th International Conference on Artificial Neural Networks*, pp. 442–447. IET, 1995.

## 7. Supplement

For all experiments data was preprocessed in the same way: 80% of normal data was used for training purposes, while remaining 20% of normal data, and 100% of abnormal data was used for testing purposes. We subsample the data by the factor of 10, and extract middle 100 time points from the whole signal. We calculate the mean and standard deviation of the training data, and use it to normalize all three datasets.

For the flow-based models, we used the following hyper-parameters: Adam optimizer with learning rate 0.01 and weight decay 0.001, ODE solver (for FFJORD+MADE) ‘dopri5’, batch size 100, number of epochs 620 (Figure 3, FFJORD+MADE), 2000 (Figure 3, MAF), 140 (Figure 3, FFJORD+MADE), 100 (Figure 3, Figure 4, MAF), 600 (Figure 4, FFJORD+MADE). Our FFJORD+MADE implementation is based on the `ffjord` library (<https://github.com/rtqichen/ffjord>) and the MADE implementation in the `pyro` library (Bingham et al., 2018). Our MAF implementation is based on a publicly available implementation (<https://github.com/ikostrikov/pytorch-flows>) using `pytorch` (Paszke et al., 2017).

For Local Outlier Factor, parameter `MinPts` is set to 50, and Chebyshev distance metric was used to calculate distance between data points. We used the implementation in `scikit-learn` (Pedregosa et al., 2011).