



energetic

D3.1

AI model for heat monitoring



Funded by the
European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them.

Contract No. 101103667

D3.1	Work Package No.	3	Task/s No.	3.1
Work Package Title	Data models Li-ion batteries			
Linked Task/s Title	Task 3.1 Battery cells thermal images Analysis			
Status	Final	(Draft/Draft Final/Final)		
Dissemination level	PU			
Due date deliverable	2024-05-31	Submission date		2024-06-10
Deliverable version	AI model for heat monitoring, 1.3			

DOCUMENT CONTRIBUTORS

Deliverable responsible	SnT		
Contributors	Organization	Reviewers	Organization
Abd El Rahman SHABAYEK	SnT	Sadok Ben Yahia	TalTech
Taz AMIETSZAJEW	CU	Etienne MAISONNY	Capgemini
		Djamila Aouada	SnT
		Arunkumar RATHINAM	SnT

DOCUMENT HISTORY

Version	Date	Comment
1.0	2024-04-04	First draft completed.
1.1	2024-04-16	SnT reviews integrated.
1.2	2024-05-17	Experiments updated.
1.3	2024-06-04	Internal consortium reviews integrated.

TABLE OF CONTENTS

DOCUMENT CONTRIBUTORS	2
DOCUMENT HISTORY	2
TABLE OF CONTENTS.....	3
1. Executive Summary.....	5
<i>Data-driven Anomaly Detection Techniques</i>	5
<i>Addressing Limitations and Proposed Approach</i>	5
<i>Conclusion</i>	6
2. Introduction	7
2.1. Anomalies in EV batteries	7
2.2. Anomaly Detection in Thermal Images.....	7
2.3. Types of Anomalies in thermal images	8
2.4. Early detection of Anomalies	8
3. Anomaly detection techniques	9
3.1. Reconstruction-based techniques	10
3.2. Embedding-based techniques	10
3.3. Synthesizing-based techniques	11
3.4. Addressing existing limitations	12
4. Unsupervised Thermal based Anomaly Detection	13
4.1. Proposed model	14
4.1.1. Feature extraction.....	14
4.1.2. Domain adaptation	14
4.1.3. Pseudo anomalies generation	15
4.1.4. Features Discrimination.....	15
4.2. Capturing thermal images	15
4.3. Thermal anomalies.....	17

4.4. Experiments	18
5. Conclusion.....	20
References	22
Appendix	25
Examples of normal heat propagation	25
Examples of overheating	26
Examples of non-homogenous heat propagation.....	26
Examples of overheating and non-homogenous heat propagation	27

1. Executive Summary

This report addresses the critical issue of ensuring safety and reliability in Electric Vehicles (EVs) by focusing on early anomaly detection in Lithium-ion batteries. These batteries are susceptible to various anomalies, with thermal deviations being a major concern. Our approach leverages thermal imaging to identify such anomalies while coping with limitations like computational constraints, limited training data, noise, and dynamic operating conditions.

Data-driven Anomaly Detection Techniques

Data-driven techniques leverage machine learning and deep learning for anomaly pattern extraction from historical battery data. Although our model is data driven, it is agnostic to the specific battery model.

There are three main categories within data-driven techniques:

1. **Reconstruction-based:** Analyze deviations from reconstructed data points. While effective, these techniques can struggle with anomalies that share features with normal data.
2. **Embedding-based:** Project data points into a lower-dimensional space for easier anomaly identification. The main limitations are the domain mismatch of the embedded features and the computational cost.
3. **Synthesizing-based:** Focuses on learning the distribution of normal data and flagging outliers. We leverage this concept for anomaly synthesis in the feature space.

Addressing Limitations and Proposed Approach

The proposed approach **FAUAD** (Feature-Adapted Unsupervised Anomaly Detection) tackles the limitations in existing techniques to effectively detect anomalies in battery thermal images thanks to:

- **Feature space anomaly synthesis** within the embedded feature space for efficient and accurate anomaly representation.
- **Domain adaptation** to adjust pre-trained features for improved effectiveness on battery thermal images.

- **Efficient architecture** that is simple but efficient for faster training, inference, and deployment in real-world applications.

Benefits:

- **Reduced training data needs:** Eliminates the need for extensive real anomaly data through unsupervised learning.
- **Universal applicability:** Adaptable to diverse and unexpected anomaly types, even those not explicitly included in training data.

Current Limitation:

Pinpointing anomaly cause: While anomaly detection is achieved, identifying the specific cause requires dedicated models with labeled data (not available due to anomalous data scarcity).

Conclusion

This report proposes FAUAD, a novel unsupervised anomaly detection method for battery thermal images in Electric Vehicles (EVs). Leveraging readily available normal data (due to anomalous data scarcity), FAUAD tackles the critical need for early detection of thermal anomalies which represent a major threat to battery health. Experiments show that FAUAD outperforms existing techniques, achieving the highest anomaly detection accuracy while maintaining a compact 15 MB model size ideal for resource-constrained edge devices in EVs. This paves the way for improved EV battery safety and reliability.

2. Introduction

Electric vehicles (EVs) are rapidly gaining traction as a sustainable transportation solution. However, ensuring their safety and reliability hinges on the health and performance of their lithium-ion batteries. These batteries are complex systems susceptible to various anomalies, from degradation and overheating to internal short circuits. Early detection of such anomalies is crucial for preventing catastrophic failures, safeguarding passengers, and extending battery life.

2.1. Anomalies in EV batteries

Anomalies in EV batteries can manifest in various ways [Berg 2015, Castelvechi 2021, Dong and Lin, 2021, Zhang et al., 2023]:

- **Degradation:** Gradual loss of battery capacity over time due to chemical and physical processes, reducing range and performance.
- **Voltage fluctuations:** Abnormal voltage spikes or dips, potentially indicating internal cell imbalances, short circuits, or faulty connections.
- **Current imbalances:** Unequal current distribution among battery cells, leading to premature aging and potential cell damage.
- **Impedance changes:** Increases in internal resistance, reflecting degradation, structural issues, or manufacturing defects.
- **Thermal deviations:** Excessive heating or uneven temperature distribution, suggesting improper cooling, internal resistances, or thermal runaway risks.

This deliverable, which addresses heat monitoring, primarily concentrates on **thermal deviations** as the key indicators of anomalies.

2.2. Anomaly Detection in Thermal Images

Thermal imaging captures the infrared radiation emitted by objects, providing valuable information about their temperature distribution. Thermal images can manifest deviations from normal thermal patterns, indicating potential problems or suspicious activities. Identifying these deviations (anomalies) effectively is crucial for making informed decisions based on thermal data.

Several factors **complicate anomaly detection in thermal images**:

- **Computational constraints:** Real-time anomaly detection in resource-limited onboard systems necessitates efficient and lightweight algorithms.
- **Limited training data:** Collecting large, diverse datasets for training anomaly detection models can be challenging. It's important to understand that creating a complete collection of anomalies to cover every case is virtually unattainable, particularly when it comes to the varied and unpredictable irregularities found in batteries.
- **Noise:** Sensor noise, environmental factors, and image acquisition conditions can introduce artifacts and obscure anomalies.
- **Dynamic operating conditions:** Battery behaviour heavily depends on driving patterns, ambient temperature, and charging regimes, necessitating adaptable anomaly detection algorithms.

The proposed approach aims **to cope with these limitations** by providing a lightweight model that eliminates the need for extensive real anomaly data to be universal enough to adapt to unexpected anomalies.

2.3. Types of Anomalies in thermal images

Anomalies in thermal images can broadly be categorized as:

- **Temperature outliers:** Significant deviations from the expected temperature range, such as hot spots indicating equipment overheating or cold spots suggesting insulation breaches.
- **Texture variations:** Irregularities in the spatial distribution of temperature, potentially revealing material defects, surface anomalies, or foreign objects.
- **Shape abnormalities:** Deviations from the expected shape of an object, indicating structural damage, equipment malfunctions, or suspicious activities.

2.4. Early detection of Anomalies

Early detection refers to the process of identifying these anomalies at the earliest possible stage, often before they have caused significant damage or disruption. This is achieved by continuously monitoring the thermal images and detecting any changes beyond what has been defined as a normal image. Any significant change in the temperature outliers, texture variations, or shape abnormalities is

considered an anomaly and triggers an alert for further investigation. This approach offers significant benefits for early detection of possible battery failures, potentially preventing catastrophic events.

3. Anomaly detection techniques

There are various approaches employed for EV battery anomaly detection:

- **Model-based techniques:** Utilize physical and electrochemical models of battery behaviour to identify deviations from predicted responses.
- **Data-driven techniques:** Leverage machine learning and deep learning algorithms to analyse historical battery data and extract anomaly patterns.
- **Hybrid techniques:** Combine the strengths of model-based and data-driven techniques for increased accuracy and interpretability.

The goal is to get a model that is **agnostic to the battery model**; hence, **data-driven techniques** will be employed. Hence, the approach can **be universally applicable** to diverse and unexpected anomaly types, as they cannot be explicitly captured for training.

There are three main categories of **data-driven anomaly detection techniques**:

1. **Reconstruction-based techniques:** try to reconstruct normal data points and identify anything that deviates significantly from the reconstruction as an anomaly.
2. **Embedding-based techniques:** map data points into a lower-dimensional space, where anomalies are expected to be isolated or deviate from the typical patterns of normal data.
3. **Synthesizing-based techniques:** aim to learn the underlying distribution of normal data and then flag any data points not well-represented by the learned model as anomalies.

3.1. Reconstruction-based techniques

These techniques exploit the notion that anomalous image regions deviate significantly from the patterns observed in training data, making their faithful reconstruction challenging. This approach leverages various techniques:

- **Generative models:** Some techniques employ generative models like autoencoders [Gong et al., 2019] and generative adversarial networks (GANs) [Goodfellow et al., 2014] to learn a compressed representation and reconstruct normal data from it. Deviations from this learned representation are then flagged as anomalies.
- **Inpainting:** Other techniques [Haselmann et al., 2018, Ristea et al., 2022, Zavrtanik et al., 2021b] frame anomaly detection as an inpainting problem. Here, random image patches are masked, and neural networks are trained to predict and fill the missing information. The structural similarity index (SSIM) loss function [Wang et al., 2004] is commonly used during training to guide the reconstruction process towards preserving structural details.
- **Anomaly map generation:** Finally, an anomaly map can be generated by calculating the pixel-wise difference between the original image and its reconstructed version. This map highlights regions with significant reconstruction errors, potentially indicating anomalies.

However, if anomalies share common features like local edges with normal data, or the reconstruction model is overly powerful, accurate reconstruction of anomalies might occur, leading to false negatives [Zavrtanik et al., 2021b].

3.2. Embedding-based techniques

These techniques have recently emerged as a powerful approach for anomaly detection. They work by compressing normal features. The features are extracted and embedded into a lower-dimensional space. This compression allows for easier identification of anomalies, which typically lie far away from the clusters formed by normal features.

Many techniques [Defard et al., 2021, Deng and Li, 2022, Roth et al., 2022, Rudolph et al., 2022] utilize pre-trained models to extract relevant features from the images. These models help improve performance by providing a good starting point for feature extraction.

However, these techniques face several limitations:

- **Domain mismatch:** real-time images often have different visual characteristics compared to the data used to pre-train models. This mismatch can lead to inaccurate anomaly detection.
- **Performance limitations:** Some techniques, like PaDiM [Defard et al., 2021] and PatchCore [Roth et al., 2022], rely on computationally expensive steps like calculating covariance inverses or searching large memory banks. This can hinder real-time performance, especially on resource-constrained devices.
- **Memory consumption:** Techniques like normalizing flow [Rezende and Mohamed, 2015], employed by methods like CS-Flow [Rudolph et al., 2022], CFLOW-AD [Gudovskiy et al., 2022], and DifferNet [Rudolph et al., 2021], can be memory-intensive as they require processing full-sized feature maps and utilizing memory-hungry layers.
- **Increased complexity:** Distillation techniques [Bergmann et al., 2020, Deng and Li, 2022] require training in two separate networks, a teacher and a student, doubling the computational cost and potentially slowing down inference.

3.3. Synthesizing-based techniques

These techniques adopt a different approach, aiming to learn how to differentiate normal data from anomalies without actual anomalous examples. This strategy involves:

- **Synthesizing anomalies on clean images:** Several techniques, like DREAM [Zavrtanik et al., 2021a], attempt to generate realistic "just-out-of-distribution" patterns and train a network to discriminate between them and normal data. However, this approach can be computationally expensive and may not capture the full spectrum of real-world anomalies.
- **CutPaste strategy** [Li et al., 2021]: This simpler technique randomly cuts and pastes image patches within an image, creating altered versions that serve as synthetic anomalies. However, these "anomalies" often lack the characteristics of real ones, limiting their effectiveness.
- **Synthesizing anomalies in feature space:** [Liu et al., 2023, Sohn et al., 2021] synthesize anomalies in the feature space by adding Gaussian noise or data augmentation. These are interesting

approaches and demonstrate high accuracy. Hence, this work is inspired by such approaches while adapting the strategy to our needs. It is worth noting that generating a comprehensive set of synthetic anomalies to encompass all possible variations is practically impossible, especially when dealing with diverse and unpredictable anomalies like those in batteries.

3.4. Addressing existing limitations

To detect anomalies in thermal images of batteries, we can overcome the limitations in the existing state of art by exploiting several key strategies:

- **Feature space anomaly synthesis:** Instead of generating synthetic anomalies directly in the thermal image domain, which can be unrealistic, we generate them within the **embedded battery feature space** to allow for more efficient and accurate anomaly representation.
- **Domain adaptation:** As pre-trained models introduce a **domain mismatch** when applied to the acquired battery thermal images, we incorporate a "**feature adapter**" to adjust pre-trained features to the battery domain to improve its effectiveness.
- **Efficient architecture:** We utilize **simplified but efficient architecture** for inference, see Section 4. This design choice facilitates faster training, inference, and deployment, making it suitable for real-applications and industrial settings.

Hence, our approach would:

- **eliminate the need for extensive real anomaly data.** Unlike traditional methods requiring large amounts of labeled anomaly data, our approach employs **unsupervised learning**. This allows the model to learn from normal data and identify deviations without pre-defined anomaly examples, making it more adaptable to unforeseen anomalies in real-world scenarios.
- **be universally applicable to diverse and unexpected anomaly types** that cannot be explicitly captured during training. This offers significant benefits for early detection of possible battery failures, potentially preventing catastrophic events.

Current limitation: While our approach would support early anomaly detection, a key limitation remains; pinpointing the specific cause of the detected anomaly. This would require dedicated models trained on labeled data containing specific types of anomalies. These are not available to the consortium as anomalous behaviors are rare and acquiring enough real anomalous images might not be achievable. Hence, the consortium adopted the unsupervised learning direction which models the normality (where enough data is available) rather than abnormality.

4. Unsupervised Thermal based Anomaly Detection

We call the proposed approach **FAUAD** (Feature-Adapted Unsupervised Anomaly Detection), see Figure 1. It models the normality of the input data and synthesizes anomalies (pseudo anomalies) in its feature space. It follows the ideology of One Class Classification (OCC) anomaly detection algorithms where a model is trained on what is considered “normal” data and then using that model to detect whether new data is normal or an anomaly [Liu et al., 2023].




Figure 1 A simplified illustration of the Energetic model called FAUAD for unsupervised anomaly detection in thermal images of batteries.

Given an input thermal image:

1. Features are extracted using the pretrained model to enrich the feature space.
2. The features pass a domain adaptation step.
3. The adapted features are then disturbed using noise to act as pseudo anomalies.
4. The adapted and disturbed features pass through a discriminator.

The above procedure has proven its effectiveness as will be seen in the reported experiments in Section 4.4.

4.1. Proposed model

As briefly explained above, the proposed algorithm extracts normal features, adapt them, produce pseudo anomalies, and finally discriminate between them. The following subsections will give more details on each step.

4.1.1. Feature extraction

The Feature Extractor, following the approach described in [Roth et al., 2022, Liu et al., 2023], obtains local features. For each image, the pre-trained network extracts features from various hierarchies. However, since the pre-trained network is biased towards the dataset on which it was trained, it is rational to select only a subset of early levels that contain more generic features.

To ensure compatibility with edge devices and enable real-time applications with EVs (Electric Vehicles), we opted for a lightweight pretrained [resnet18](#) model in PyTorch [PyTorch, 2024] where the features from the first three layers are used, similar to [Akçay et al., 2022]. By selecting this model, we ensure that our solution can be easily deployed on various edge devices without compromising performance.

4.1.2. Domain adaptation

To address the difference in distribution between the captured thermal images and the dataset used for backbone pre-training, we utilize a Feature Adaptor. This Adaptor is employed to transfer the training features to the target domain, ensuring compatibility between the two domains. the adaptor component

performs a transformation on the pre-trained features. In our case, the feature adaptor consists of a single fully connected layer that has the same number of input and output channels, and it does not include a bias term [Liu et al., 2023].

4.1.3. Pseudo anomalies generation

Training the Discriminator to identify normal samples requires negative samples, or examples of anomalies. While sampling real-world thermal anomalies can be difficult, adding Poisson noise to normal features provides a tractable way to generate artificial anomalies [Altmann et al., 2021]. Poisson noise is a natural choice because it models the occurrence of rare events, which is a common characteristic of anomalies in real-world data.

4.1.4. Features Discrimination

Acting as a normality scorer, the Discriminator learns to distinguish between normal and anomalous images. During training, we present it with both normal data (positive samples) and data augmented with Poisson noise (negative samples). The Discriminator's goal is to output high values for genuine data (normal features) and low values for anomalies (features with noise). We achieve this with a simple two-layer multi-layer perceptron (MLP), a common architecture for classification tasks.

4.2. Capturing thermal images

For AI thermal image analysis model training, real thermal images had to be obtained. This has been performed by Coventry University, using FLIR T640 infrared camera by FLIR Systems. This system is rated at +/-2% of reading accuracy, capturing 640 x 480 pixel images at 24 bit depth using sRGB colour representation. Focal length of 13mm and exposure time of 1/46 sec. are pre-set.

The camera is suspended above the cell under testing (SPIM11309102-GL40 Pouch cell), as shown in Figure 2. The cell under test is connected to a 5V 200A 4-wire battery cycler (Neware®). All testing is conducted inside a controlled environment thermal chamber (Weiss Technik®).

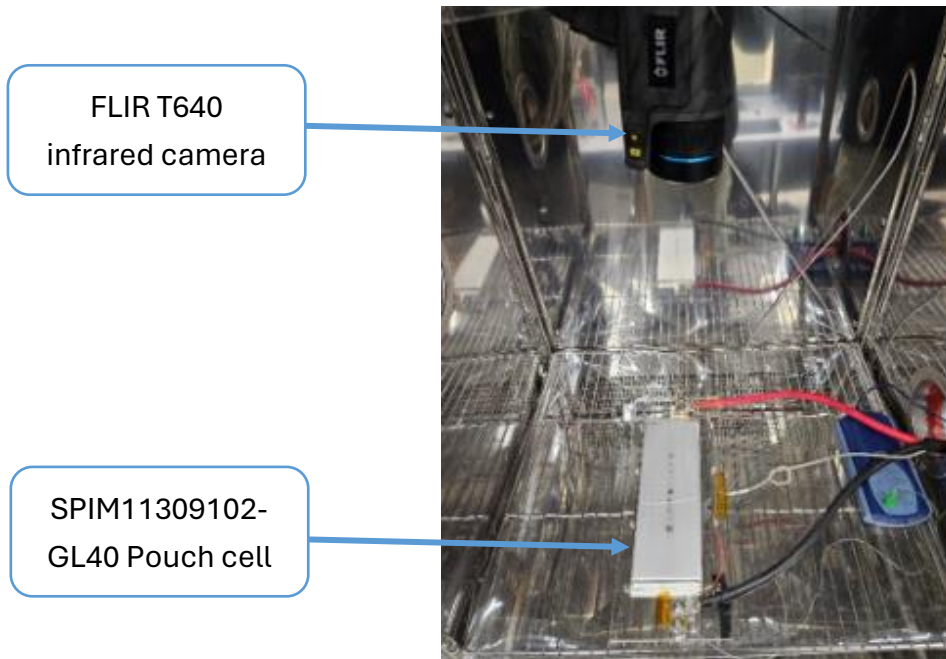


Figure 2 Thermal images acquisition test set-up

Initially obtained images contained reflections, which leads to erroneous thermal readings, despite attempts to keep the chamber blacked out during tests. To alleviate this issue the surface of the cell was sprayed with a matt black paint (Ambersil) as advised by the thermal camera manufacturer, see Figure 3. Following this correction, Example of the resulting thermal images is shown in Figure 4.



Figure 3 Cell sprayed with matt black paint (Ambersil).

Thermal images were collected using a time-lapse mode of 15 seconds per image, resulting in thousands of images per test. Cell cycling performance data (potential, current, capacity) was collected alongside. The combined data obtained was shared with SnT for AI image analysis training.




Figure 4 Thermal image of a cell captured during cycling. Left) before painting, Right) after painting.

4.3. Thermal anomalies

Thermal anomalies can be defined by delta over space / gradient (e.g. one area is hotter by 5 degrees - normal, by 20 degrees - anomalous), and time (steady increase over time - normal, sudden increase – anomalous), so they exist across both spatial and temporal axes. In thermal images, “overheat” refers to areas that display colors towards the higher end of the used colormap, indicating temperatures that exceed the normal or safe range for the observed object or environment.

The proposed unsupervised anomaly detection model FAUAD is tested and benchmarked against state-of-the-art models. Two sets of anomalies were extracted from the captured thermal images: overheat images and the images with non-uniform thermal propagation (represented by reflections), see Figures 5, 6.




Figure 5 Thermal image of a cell captured during cycling. Left) overheated, Right) with reflections before matte black painting.

The first category represents dangerous levels of heat on the battery surface and the second category visually represents non-homogeneous abnormal heat propagation. Reflections observed during methodology development, as outlined in section 4.2, have been utilized to replicate non-homogeneity. While coating the cell with matte black paint, an area was masked and then uncovered, leaving a reflective patch, see Figure 6. That patch resulted in spatial anomaly, leaving

areas of sharp temperature gradient, which would indicate abnormal cell operation and areas of potential cell failure.



Figure 6 Coating the cell with matte black paint, an area was masked and then uncovered, leaving a reflective patch. Left) the painted cell, Right) the thermal image.

The captured thermal images have been structured to extend the MVTec [Bergmann et al., 2019] industrial anomaly detection dataset for the classification task. The “battery anomaly detection dataset” folder contains “train” and “test” folders. The train folder contains normal data only in a folder called “good”. The “test” folder contains the “good” folder for normal data and remaining folders (“overheat” and “reflection”) to represent anomalies. For more examples, see Appendix A.

It is important to note that the model learns what is ‘normal’ based on the input data characterized as ‘normal.’ Consequently, any pattern that deviates from this learned normality, when introduced to the trained model, will be flagged as an anomaly. For instance, consider a scenario where a thermal camera is dislodged from its original position due to a significant jolt. This displacement may result in the production of thermal images with propagation patterns that differ markedly from the ‘normal’ patterns on which the model was trained. In such a case, the model would identify the observed images as anomalies. Although these are not genuine anomalies from the perspective of the battery’s functionality, the triggered anomaly alert would still indicate that an abnormal event has occurred.

4.4. Experiments

We compared FAUAD’s performance to state-of-the-art algorithms using two classification scenarios:

- The first scenario mimicked real-world data by including a small percentage of anomalies (10%), simulating potential issues during data collection. Although this might not be a realistic scenario in the context of batteries as the environment will be well controlled for safety reasons, it provides an opportunity to assess the effectiveness of FAUAD in learning normality compared to existing studies.

- The second experiment acted as a control, training on only clean, normal data.

This two-pronged approach allows us to assess the robustness of all models under realistic and ideal data collection conditions.

For feature extraction in all models, we opted for a lightweight backbone, pre-trained ResNet18 model from PyTorch's TorchVision library. ResNet18's pre-trained weights, obtained from its training on the massive ImageNet dataset, provide a valuable foundation for anomaly detection. These pre-trained features capture generic image properties that can be leveraged to model normality in our task. This choice strikes a balance between model size and performance, making it ideal for deployment on resource-constrained edge devices.

Table 1 Experimental scenarios: First scenario (first column) mimicked real-world data simulating potential issues during data collection. Second scenario (second column) to train on normal data only.

Technique	AUROC (normal data contaminated with 10% anomalies)	AUROC (clean normal data)	Trained model size
FAUAD (Energetic's model)	0.990	1	15 MB
SimpleNet [Liu et al., 2023]	0.977	1	15 MB
DRAEM [Zavrtanik et al., 2021a]	0.922	0.991	1100 MB
CFA [Lee et al., 2022]	0.878	0.942	30MB
STFPM [Wang et al., 2021]	0.871	0.961	43 MB
PaDiM [Defard et al., 2021]	0.863	0.996	175 MB
EfficientAD [Batzner et al., 2024]	0.810	1	81 MB
DFM [Nilesh et al., 2019]	0.803	0.996	15 MB
FastFlow [Yu et al., 2021]	0.793	1	65 MB
PatchCore [Roth et al., 2022]	0.774	0.990	42 MB
CFLOW-AD [Gudovskiy et al., 2022]	0.769	0.873	171 MB

This table compares FAUAD with the state-of-art anomaly detection techniques based on Area Under the ROC Curve (AUROC), model size, and performance on clean vs. contaminated data. Here are some key takeaways:

- **FAUAD (Energetic's model):** This model boasts the highest AUROC (0.990) on contaminated data and a perfect score (1) on clean data. Additionally, its model size is a modest 15 MB. This combination makes it a very attractive option for edge devices.

- **SimpleNet:** While not quite reaching the heights of FAUAD, SimpleNet delivers a strong AUROC (0.977) on contaminated data and maintains a perfect score on clean data. It also shares the same 15 MB model size as FAUAD, making it another strong contender.
- **Performance on Clean Data:** Techniques like DFM (15 MB), CFA (30 MB), PatchCore (42 MB), STFPM (43 MB), FastFlow (65 MB) and EfficientAD (81 MB), have lower overall performance on contaminated data. However, they excel at identifying normal data (AUROC is equal or close to 1).

In conclusion, FAUAD and SimpleNet emerge as the top contenders based on their high AUROC on contaminated data, perfect score on clean data, and compact model size.

5. Conclusion

This report addressed the crucial challenge of ensuring safety and reliability in Electric Vehicles (EVs) by proposing FAUAD, a novel approach for unsupervised anomaly detection in Lithium-ion battery thermal images. Thermal anomalies pose a significant threat to battery health, and FAUAD offers an efficient and robust solution for their early detection to help mitigate any significant threat due to unexpected failures.

A key strength of FAUAD in battery health monitoring lies in its ability to leverage unsupervised learning. This approach is particularly well-suited due to the scarcity of labeled thermal image-based anomaly data. By focusing on readily available normal data, FAUAD can effectively detect anomalies without the need for extensive and potentially hazardous anomaly collection.

The conducted experiments demonstrate FAUAD's effectiveness. Amongst the evaluated techniques, FAUAD and SimpleNet emerge as the frontrunners. Both achieve a perfect AUROC score on clean data, indicating exceptional ability to identify normal battery behavior. Furthermore, FAUAD boasts the highest AUROC (0.990) on normal data contaminated with few anomalies, highlighting its proficiency in anomaly detection even under less-than-ideal conditions. This is achieved while maintaining a compact model size (15 MB), making it suitable for deployment on resource-constrained edge devices commonly found in EVs.

In conclusion, FAUAD presents a significant advancement in unsupervised anomaly detection for Lithium-ion batteries. Its ability to leverage readily available normal data, combined with its strong performance and efficient design, makes FAUAD a compelling solution for ensuring EV battery safety and reliability. Future work could explore incorporating labeled data, if it becomes available, to potentially improve FAUAD's ability to pinpoint the root cause of detected anomalies.

References

[Akçay et al., 2022] Samet Akçay, Dick Ameln, Ashwin Vaidya, Barath Lakshmanan, Nilesh Ahuja, Utku Genc. Anomalib: A deep learning library for anomaly detection, IEEE International Conference on Image Processing (ICIP), pages 1706-1710, 2022.

[Altmann et al., 2021] Yoann Altmann, Dan Yao, Stephen McLaughlin, Mike E. Davies. Robust Linear Regression and Anomaly Detection in the Presence of Poisson Noise Using Expectation-Propagation. In: Gelman, L., Martin, N., Malcolm, A.A., (Edmund) Liew, C.K. (eds) Advances in Condition Monitoring and Structural Health Monitoring. Lecture Notes in Mechanical Engineering. Springer, 2021.

[Batzner et al., 2024] Kilian Batzner, Lars Heckler, Rebecca König: EfficientAD: Accurate Visual Anomaly Detection at Millisecond-Level Latencies. arXiv:2303.14535v3, 2024.

[Berg 2015] Berg H. Batteries for Electric Vehicles: Materials and Electrochemistry. Cambridge University Press; 2015.

[Bergmann et al., 2019] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9592–9600, 2019.

[Bergmann et al., 2020] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4183–4192, 2020.

[Castelvecchi 2021] Castelvecchi, D. Electric cars: the battery challenge. Nature 596, 336–339, 2021.

[Defard et al., 2021] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. PaDiM: a patch distribution modeling framework for anomaly detection and localization. In International Conference on Pattern Recognition, pages 475–489. Springer, 2021.

[Deng and Li, 2022] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022.

[Dong and Lin, 2021] Guangzhong Dong and Mingqiang Lin. Model-based thermal anomaly detection for lithium-ion batteries using multiple-model residual generation. In Journal of Energy Storage, Volume 40, 2021.

[Gong et al., 2019] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In Proceedings of IEEE/CVF International Conference on Computer Vision, pages 1705–1714, 2019.

[Goodfellow et al., 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems, 27, 2014.

[Gudovskiy et al., 2022] Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. CFLOW-AD: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 98–107, 2022.

[Haelmann et al., 2018] Matthias Haselmann, Dieter P Gruber, and ul Tabatabai. Anomaly detection using deep learning based image completion. In 2018 17th IEEE international conference on machine learning and applications (ICMLA), pages 1237–1242. IEEE, 2018.

[Lee et al., 2022] Sungwook Lee, Seunghyun Lee and Byung Cheol Song: CFA: Coupled-Hypersphere-Based Feature Adaptation for Target-Oriented Anomaly Localization. IEEE Access 10: 78446-78454, 2022.

[Li et al., 2021] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 9664–9674, 2021.

[Liu et al., 2023] Zhikang Liu, Yiming Zhou, Yuansheng Xu, and Zilei Wang. SimpleNet: A Simple Network for Image Anomaly Detection and Localization. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 20402- 20411, 2023.

[Nilesh et al., 2019] Nilesh .A.Ahuja, Ibrahima J. Ndiour, Trushant Kalyanpur and Omesh Tickoo: Probabilistic Modeling of Deep Features for Out-of-Distribution and Adversarial Detection. In: Bayesian Deep Learning Workshop (2019).

[PyTorch, 2024], 20 March 2024,
<https://pytorch.org/vision/main/models/generated/torchvision.models.resnet18.html>.

[Rezende and Mohamed, 2015] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In International conference on machine learning, pages 1530–1538, 2015.

[Ristea et al., 2022] Nicolae-Catalin Ristea, Neelu Madan, Radu Tudor Ionescu, Kamal Nasrollahi, Fahad Shahbaz Khan, Thomas B Moeslund, and Mubarak Shah. Self-

supervised predictive convolutional attentive block for anomaly detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 13576–13586, 2022.

[Roth et al., 2022] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Scholkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14318–14328, 2022.

[Rudolph et al., 2021] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but different: Semi-supervised defect detection with normalizing flows. In Proceedings of IEEE/CVF winter conference on applications of computer vision, pages 1907–1916, 2021.

[Rudolph et al., 2022] Marco Rudolph, Tom Wehrbein, Bodo Rosenhahn, and Bastian Wandt. Fully convolutional cross-scale-flows for image-based defect detection. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 1088–1097, 2022.

[Sohn et al., 2021] Kihyuk Sohn, Chun-Liang Li, Jinsung Yoon, Minho Jin, and Tomas Pfister: Learning and evaluating representations for deep one-class classification. In: International Conference on Learning Representations (2021).

[Wang et al., 2004] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing, 13(4):600–612, 2004.

[Wang et al., 2021] Guodong Wang, Shumin Han, Errui Ding and Di Huang: Student-Teacher Feature Pyramid Matching for Unsupervised Anomaly Detection. In: British Machine Vision Conference (2021).

[Yu et al., 2021] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao and Liwei Wu: FastFlow: Unsupervised Anomaly Detection and Localization via 2D Normalizing Flows. CoRR abs/2111.07677 (2021).

[Zavrtanik et al., 2021a] Vitjan Zavrtanik, Matej Kristan, and Danijel Skocaj. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. Proceedings IEEE International Conference on Computer Vision, pages 8330–8339, 2021.

[Zavrtanik et al., 2021b] Vitjan Zavrtanik, Matej Kristan, and Danijel Skocaj. Reconstruction by inpainting for visual anomaly detection. Pattern Recognition, 112:107706, 2021.

[Zhang et al., 2023] Zhang, J., Wang, Y., Jiang, B. et al. Realistic fault detection of li-ion battery via dynamical deep learning. Nat Commun 14, 5940 (2023).

Appendix A

The following sections present thermal images captured using the method described in Section 4.2. Note that the ambient room temperature during capture was approximately 24°C.


The maximum temperature displayed in the images may vary depending on the thermal camera settings. There are two main options:

- **Automatic:** In this mode, the colormap is relative and doesn't directly translate to normal vs. abnormal behavior. It simply shows a temperature gradient from coldest (typically dark blue) to hottest (typically red or yellow).
- **Fixed:** This mode sets a specific temperature range for the colormap. This allows the images to visually represent normal operation (colors within a safe range) and unsafe operation (colors exceeding a pre-defined threshold, typically above 50°C).

Using **the fixed mode of the colormap is crucial** for accurate analysis.


Examples of normal heat propagation

Under normal operation, a thermal image of a battery should display a consistent pattern of heat dissipation.




Examples of overheating


While the heat distribution appears uniform, the shades of red color indicate the presence of dangerously high temperatures.



Examples of non-homogenous heat propagation

In the absence of real non-homogenous heat propagation anomalies for testing purposes, we can leverage reflections in thermal images, see Section 4.2 and 4.3, to simulate non-uniform heat patterns. These reflections may cause colors in the image to deviate from the expected homogenous distribution on the battery surface which resembles non-homogenous propagation anomalies.





Examples of overheating and non-homogenous heat propagation

These images showcase an example where two visual anomalies coexist - uneven heat distribution across the battery surface and excessively high temperatures (overheating).

