University of Calabria
Department of Computer Engineering, Modeling, Electronics, and Systems
PhD Program in Information and Communication Technologies (ICT)
XXXVI Cycle

Machine and Deep Learning Techniques for Anomaly Detection and Generation

Thesis Advisors

Giuseppe Manco, PhD
Ettore Ritacco, PhD

Coordinator

Prof. Giancarlo Fortino

Candidate

Angelica Liguori
226126

Outline

- **Problem Definition & Motivations**
- **Personal Contribution**
  - **Anomaly Detection**
    - *A Federated-based Anomaly Detection Systems*
    - *Autoencoder-based Anomaly Detection System*
  - **Anomaly Generation**
    - *Adversarial Reconstruction Network*
  - **Ongoing**
    - *Dynamic Graph Generation for Anomaly Detection*
- **Conclusions**
- **About my Academic Career**

Problem Definition

- An **anomaly**, or **outlier**, is an observation that significantly deviates from other observations

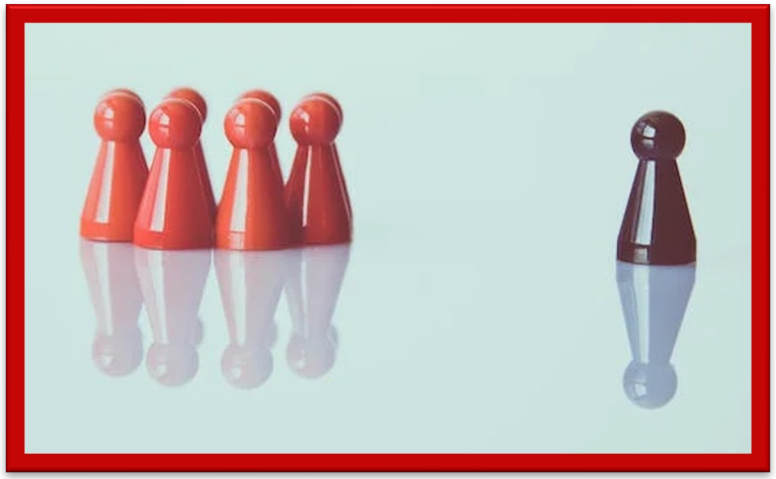

Problem Definition

- An **anomaly**, or **outlier**, is an observation that significantly deviates from other observations
- **Anomaly Detection**
  - **Goal**: Detect anomalous behaviors in a collection of data

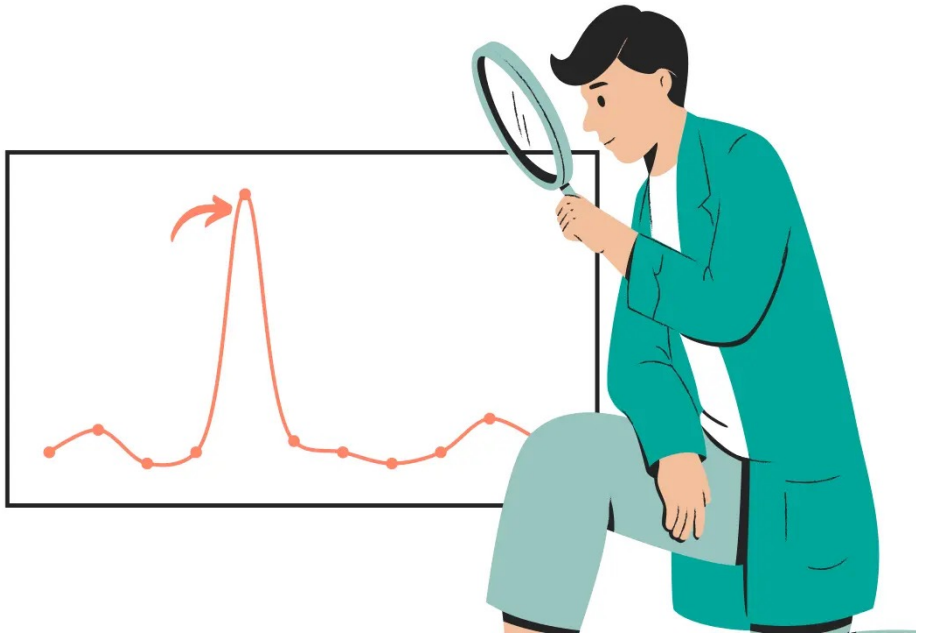

Problem Definition

- An **anomaly**, or **outlier**, is an observation that significantly deviates from other observations
- **Anomaly Detection**
  - **Goal**: Detect anomalous behaviors in a collection of data
- **Anomaly Generation**
  - **Goal**: Generate synthetic but realistic anomalies

Challenges

Scarcity of labeled anomalies

Expensive Data Collection

Labeling data is resource consuming

Data Transmission has a cost

Inherent Class Imbalance

Opportunities

Detection

Cybersecurity

Smart Industry

Fraud Analysis

Healthcare

Generation

Data enrichment

Free labeling

Class imbalance

Personal Contribution
Anomaly Detection

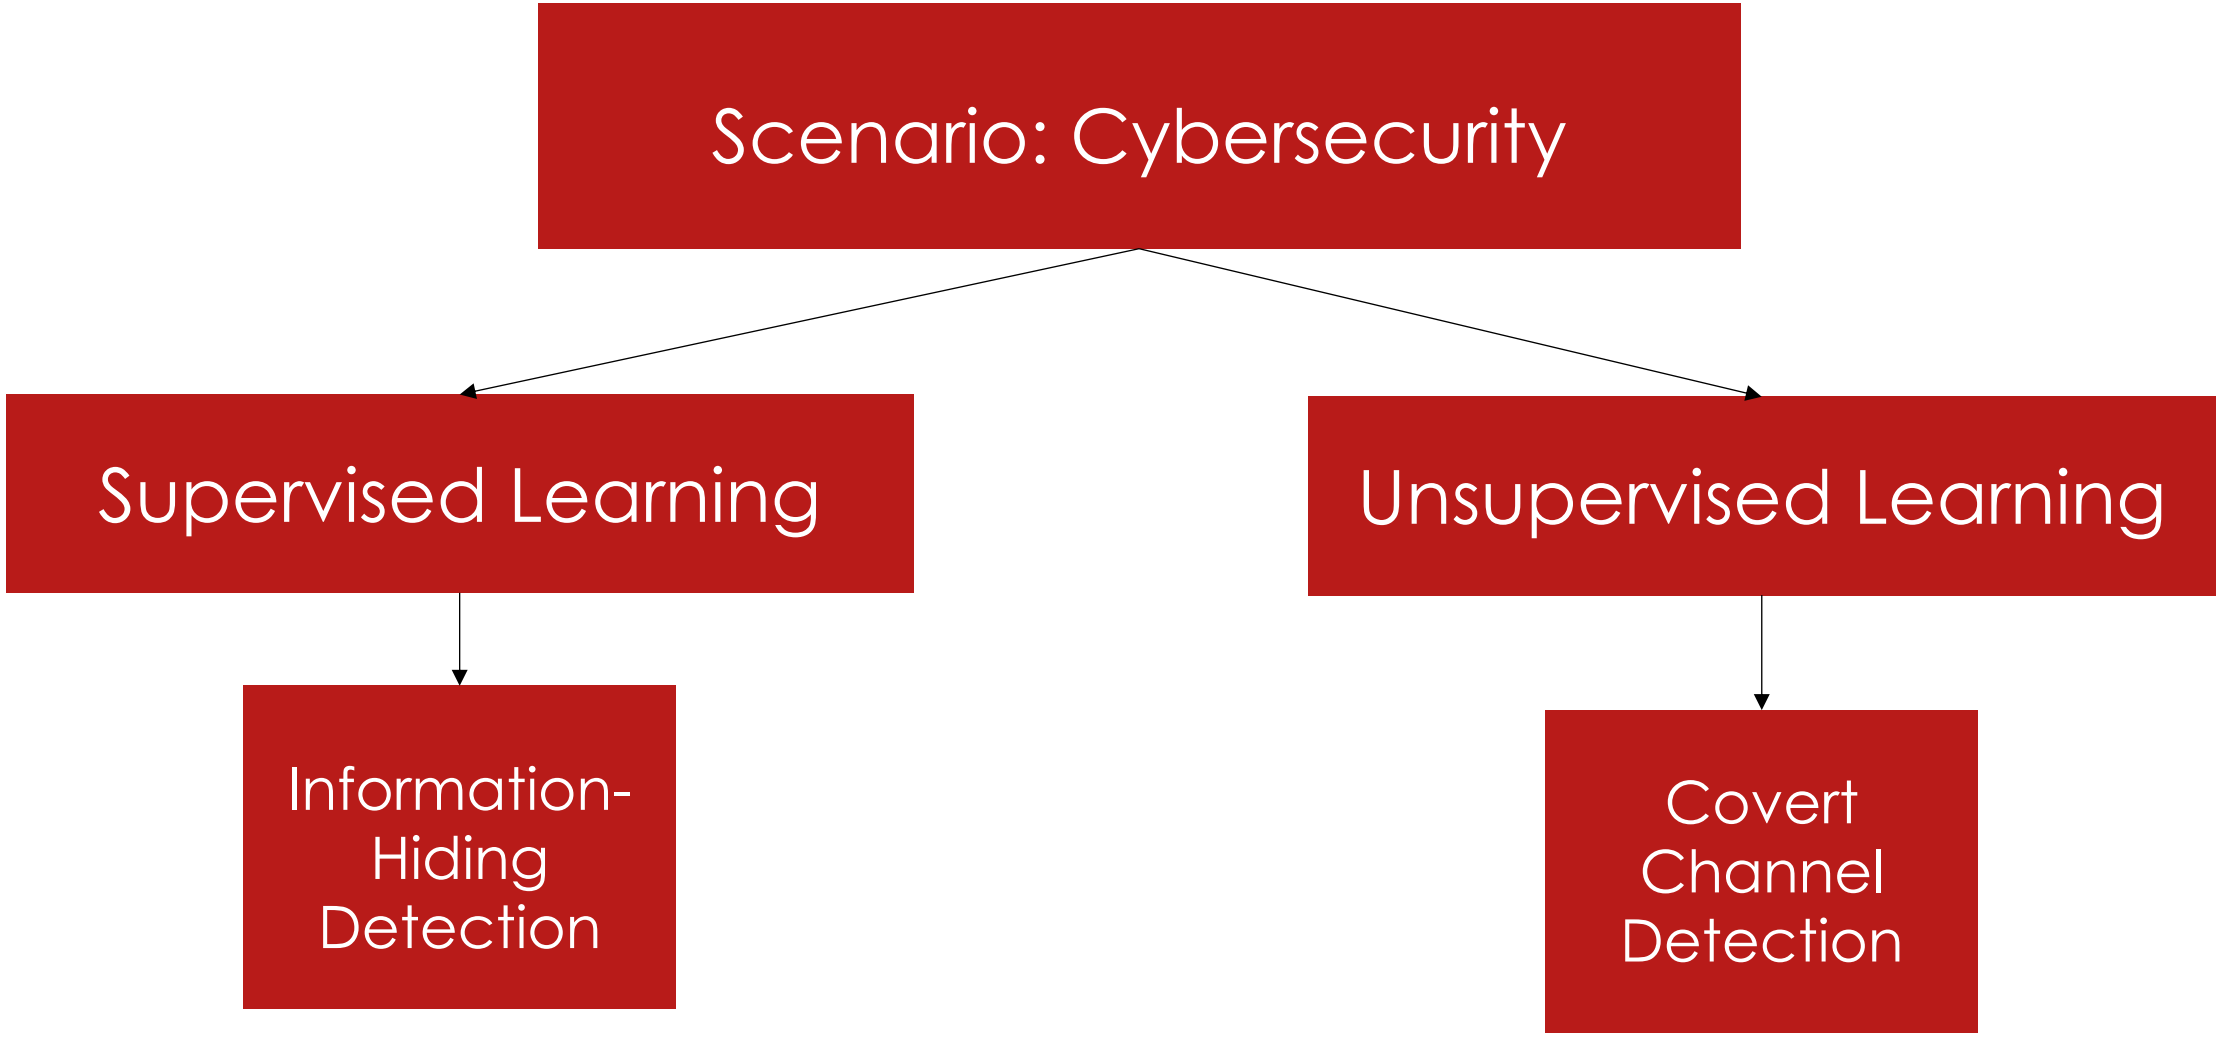

Federated-based Anomaly Detection Systems

- **Scenario**
  - A threat actor leverages steganography techniques to hide malicious payload within high-resolution icons distributed across mobile apps
- **Solution**
  - Defining a supervised federated learning approach that allows cooperation among different clients/devices to detect abnormal behaviors in the network

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

Federated-based Anomaly Detection Systems

RQ: How are the performances of the federated approach compared with a centralized one?

Approach	Coding	AUC	AUPRC	F1-Score
Centralized	Plain	**0.972**	**0.851**	**0.835**
	Base64	**0.899**	**0.605**	0.589
	zip	0.776	0.397	0.344
Federated	Plain	0.970	0.842	0.817
	Base64	0.893	0.594	**0.614**
	zip	**0.856**	**0.498**	**0.363**

FL approach achieves comparable performances with a fully centralized method w/o the necessity of moving data toward a single node.

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

Autoencoder-based Anomaly Detection Systems

- **Scenario**
  - Identify the presence of network covert channels, i.e., hidden communication, within traffic flows
- **Solution**
  - Defining an ensemble of unsupervised neural networks, i.e., encoder-decoder architectures

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Autoencoder-based Anomaly Detection Systems

RQ: Can the ensemble strategy improve the performance of the 'single' model?

Model Type	Detection Threshold	Accuracy	Precision	Recall	F1-Score
Sparse U-Net	90^{th} perc.	0.882	0.743	**0.993**	0.850
	95^{th} perc.	0.921	0.822	0.976	0.893
	99^{th} perc.	0.936	0.942	0.865	0.902
DAE	90^{th} perc.	0.869	0.724	**0.993**	0.837
	95^{th} perc.	0.910	0.801	0.975	0.880
	99^{th} perc.	0.905	**0.962**	0.750	0.843
Sparse-AE	90^{th} perc.	0.875	0.737	0.979	0.841
	95^{th} perc.	0.901	0.795	0.951	0.866
	99^{th} perc.	0.902	0.922	0.778	0.844
Skip-AE	90^{th} perc.	0.875	0.736	0.982	0.841
	95^{th} perc.	0.907	0.799	0.968	0.876
	99^{th} perc.	0.819	0.853	0.563	0.678
Ensemble (k=3)	90^{th} perc.	0.894	0.771	0.979	0.863
	95^{th} perc.	0.947	0.902	0.948	0.924
	99^{th} perc.	**0.955**	0.950	0.915	**0.932**

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Anomaly Detection & Generation

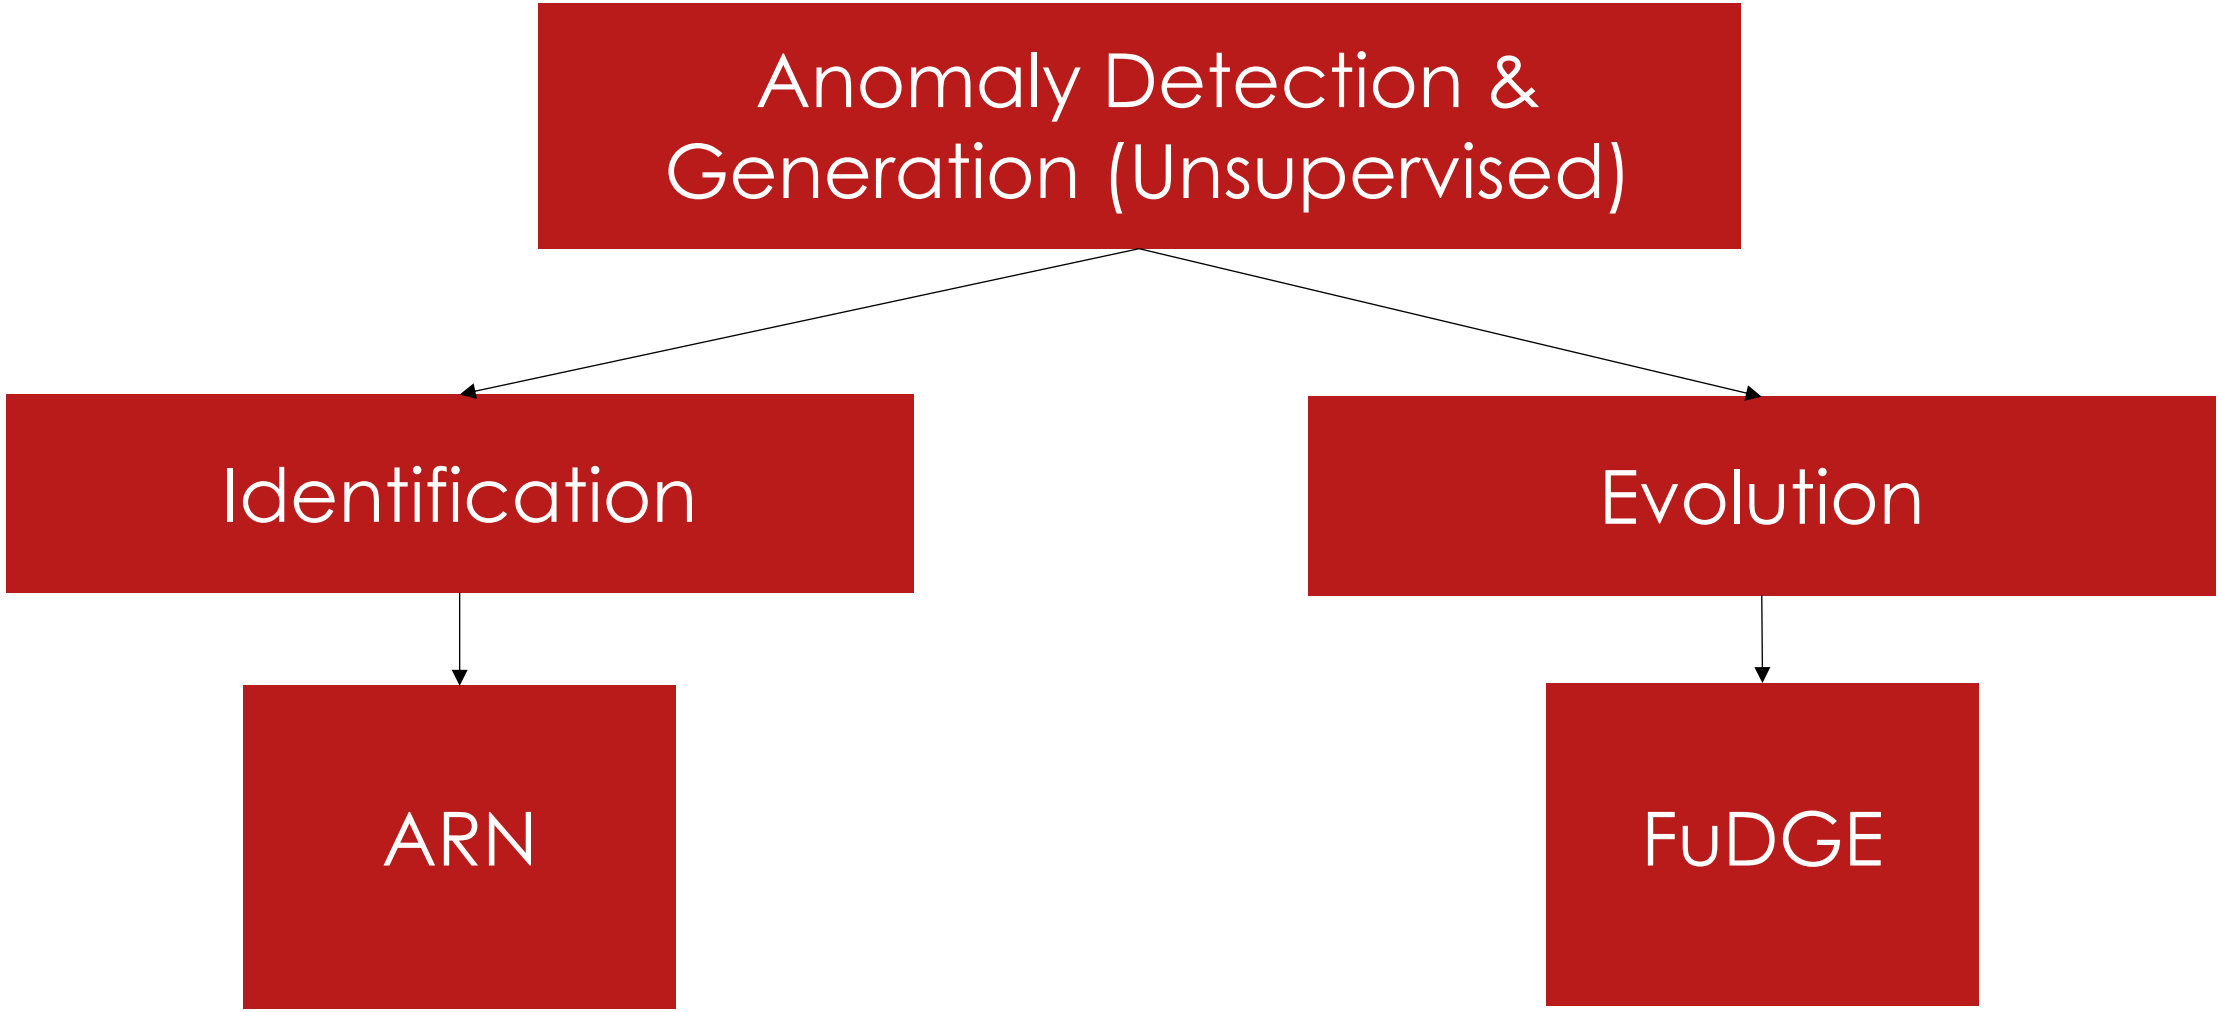

Adversarial Reconstruction Networks
Methodology

RQ: Can we build a model that effectively distinguishes between normal and abnormal behaviors using only the available normal data?

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Methodology

RQ: Can we build a model that effectively distinguishes between normal and abnormal behaviors using only the available normal data?

A: Adversarial Reconstruction Network (ARN), a twofold neural architecture aimed at generating and identifying anomalies, without needing information about them

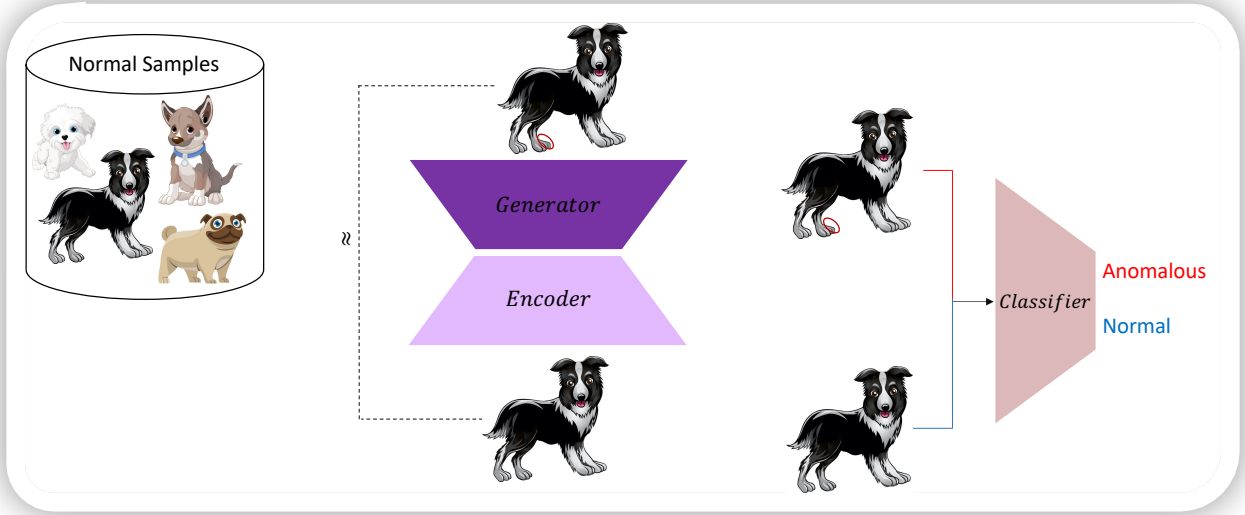

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Methodology

ARN relies on three components:

- An encoder q_ψ which can summarize the relevant features of x, i.e., an observable sample coming from a distribution $\mathbb{P}_\mathcal{D}$ ($x \sim \mathbb{P}_\mathcal{D}$), in a latent code z
- A generator g_ϕ which, given z, aims at generating a variants of x, i.e., $\tilde{x}$
- A discriminator p_θ which models the outlierness degree

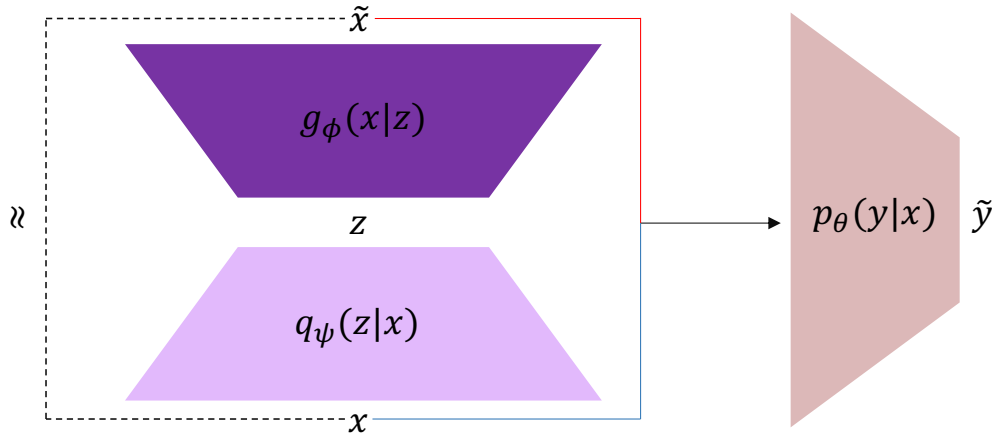

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Methodology

The adversarial game has associated the discriminator loss

$$\mathcal{L}_{\mathcal{D}}(\theta|\phi,\psi) = \mathbb{E}_{x\sim\mathbb{P}_{\mathcal{D}}}[\log p_\theta(0|x)] + \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p_\theta(1|\tilde{x})]$$

and the generator loss

$$\mathcal{L}_G(\phi,\psi|\theta) = \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p_\theta(0|\tilde{x})] + \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p(x|\tilde{x})] - \mathbb{KL}\big[q_\psi(z|x)||p(z)\big]$$

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ1. Does the outlier generator produce realistic outliers? How does it affect the predictive power?

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ1. Does the outlier generator produce realistic outliers? How does it affect the predictive power?

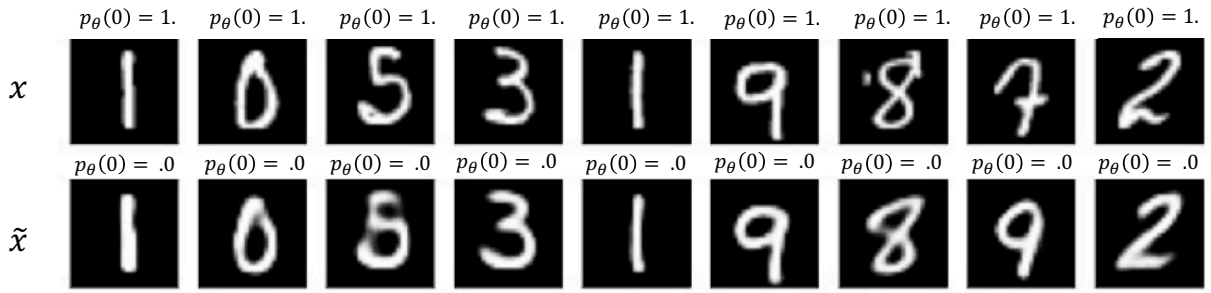

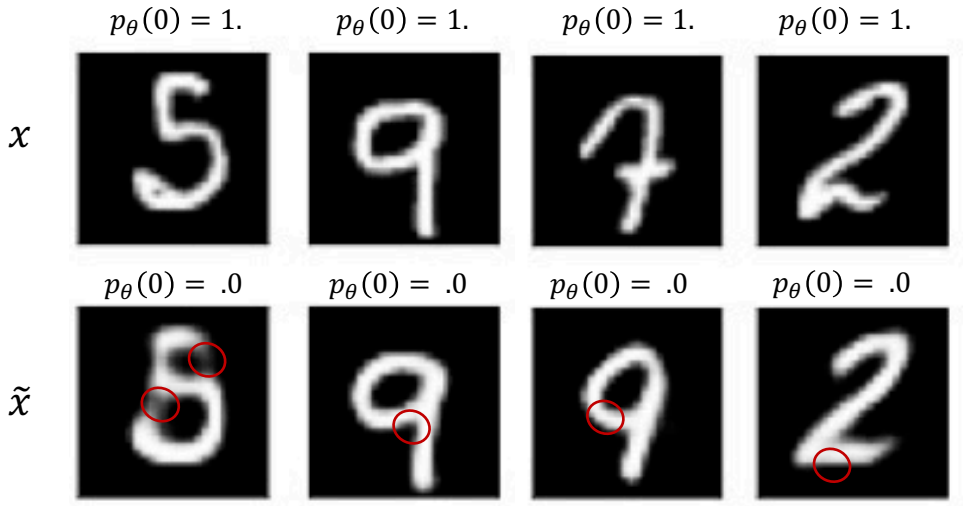

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.

- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ2. How does the predictive power of the discriminator compare to other state-of-the-art approaches?

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ2. How does the predictive power of the discriminator compare to other state-of-the-art approaches?

Dataset	ARN^G		ARN^N		FenceGAN		GANomaly		OC-SVM		Baseline	
	AUC	AUPRC	AUC	AUPRC	AUC	AUPRC	AUC	AUPRC	AUC	AUPRC	AUC	AUPRC
KDDCUP99	.99 ± .00	.99 ± .00	1.00 ± .00	.99 ± .00	.99 ± .00	.99 ± .00	1.00 ± .00	1.00 ± .00	.96 ± .00	.97 ± .00	1.00 ± .00	1.00 ± .00
KDDCUP99$_{Rev}$	.97 ± .01	.95 ± .02	**.99 ± .00**	.95 ± .02	.84 ± .01	.77 ± .01	.92 ± .01	.86 ± .01	.81 ± .00	.71 ± .00	.91 ± .01	.87 ± .01
KDDCUP99$_{Inv}$	1.00 ± .00	1.00 ± .00	1.00 ± .00	1.00 ± .00	.92 ± .03	.72 ± .08	.91 ± .04	.90 ± .03	.95 ± .00	.82 ± .00	1.00 ± .00	1.00 ± .00
NSL-KDD	.99 ± .00	.99 ± .01	.99 ± .00	.98 ± .01	.96 ± .00	.97 ± .00	.97 ± .01	.97 ± .01	.96 ± .00	.97 ± .00	.99 ± .00	.98 ± .00
DoH	.99 ± .01	1.00 ± .00	.99 ± .01	1.00 ± .00	.88 ± .02	.97 ± .00	.99 ± .00	1.00 ± .00	.88 ± .00	.97 ± .00	.96 ± .00	.99 ± .00
DoH$_{Inv}$	.98 ± .01	.97 ± .02	1.00 ± .00	**1.00 ± .00**	.89 ± .02	.44 ± .05	1.00 ± .00	.98 ± .01	.90 ± .00	.49 ± .01	.99 ± .00	.91 ± .04
CoverType	.94 ± .01	.95 ± .01	.92 ± .04	.93 ± .03	.70 ± .03	.41 ± .02	.56 ± .05	.30 ± .04	.73 ± .02	.43 ± .02	.53 ± .02	.28 ± .02
CreditCard	-	-	.99 ± .01	.59 ± .06	.90 ± .01	.51 ± .03	.84 ± .02	.36 ± .05	.92 ± .01	.57 ± .01	.99 ± .00	**.76 ± .01**
Bank	.77 ± .06	.63 ± .09	.69 ± .07	.50 ± .11	.56 ± .01	.23 ± .01	.53 ± .02	.22 ± .02	.60 ± .00	.28 ± .00	.65 ± .00	.32 ± .01

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ3. Focusing on real scenarios, how is the accuracy affected by contamination in the learning process?

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ3. Focusing on real scenarios, how is the accuracy affected by contamination in the learning process?

	KDDCUP99$_{Rev}$		KDDCUP99$_{Inv}$		CoverType	
	ARN^N	**GANomaly**	**ARN^G**	**GANomaly**	**ARN^G**	**GANomaly**
No contamination	.99 ± .00	.92 ± .01	.98 ± .02	.91 ± .04	.96 ± .01	.56 ± .05
$p = 1\%$	.98 ± .00	.88 ± .02	.98 ± .02	.88 ± .07	.70 ± .06	.54 ± .11
$p = 5\%$	.93 ± .05	.81 ± .01	.95 ± .03	.76 ± .20	.72 ± .10	.52 ± .10

	KDDCUP99	**NSL-KDD**	**DoH**
No contamination	.98 ± .01	.98 ± .01	.99 ± .01
$p = 1\%$	.96 ± .03	.98 ± .01	1.00 ± .00
$p = 5\%$	.94 ± .03	.95 ± .04	.99 ± .01
$p = 10\%$	.83 ± .09	.93 ± .03	.99 ± .00
$p = 25\%$	.74 ± .14	.91 ± .04	.98 ± .01
$p = 50\%$	.65 ± .12	.74 ± .07	.78 ± .06

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ3. Focusing on real scenarios, how is the accuracy affected by contamination in the learning process?

	KDDCUP99$_{Rev}$		KDDCUP99$_{Inv}$		CoverType	
	ARN^N	**GANomaly**	**ARN^G**	**GANomaly**	**ARN^G**	**GANomaly**
No contamination	.99 ± .00	.92 ± .01	.98 ± .02	.91 ± .04	.96 ± .01	.56 ± .05
$p = 1\%$	.98 ± .00	.88 ± .02	.98 ± .02	.88 ± .07	.70 ± .06	.54 ± .11
$p = 5\%$	.93 ± .05	.81 ± .01	.95 ± .03	.76 ± .20	.72 ± .10	.52 ± .10

	KDDCUP99	**NSL-KDD**	**DoH**
No contamination	.98 ± .01	.98 ± .01	.99 ± .01
$p = 1\%$	.96 ± .03	.98 ± .01	1.00 ± .00
$p = 5\%$	.94 ± .03	.95 ± .04	.99 ± .01
$p = 10\%$	.83 ± .09	.93 ± .03	.99 ± .00
$p = 25\%$	.74 ± .14	.91 ± .04	.98 ± .01
$p = 50\%$	.65 ± .12	.74 ± .07	.78 ± .06

In what degree a limited amount of supervision helps the learning process?

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ3. Focusing on real scenarios, how is the accuracy affected by contamination in the learning process?

	KDDCUP99_{Rev}		KDDCUP99_{Inv}		CoverType	
	ARN^N	GANomaly	ARN^G	GANomaly	ARN^G	GANomaly
No contamination	.99 ± .00	.92 ± .01	.98 ± .02	.91 ± .04	.96 ± .01	.56 ± .05
$p = 1\%$	.98 ± .00	.88 ± .02	.98 ± .02	.88 ± .07	.70 ± .06	.54 ± .11
$p = 5\%$	.93 ± .05	.81 ± .01	.95 ± .03	.76 ± .20	.72 ± .10	.52 ± .10

	KDDCUP99	NSL-KDD	DoH
No contamination	.98 ± .01	.98 ± .01	.99 ± .01
$p = 1\%$	.96 ± .03	.98 ± .01	1.00 ± .00
$p = 5\%$	.94 ± .03	.95 ± .04	.99 ± .01
$p = 10\%$	.83 ± .09	.93 ± .03	.99 ± .00
$p = 25\%$	.74 ± .14	.91 ± .04	.98 ± .01
$p = 50\%$	.65 ± .12	.74 ± .07	.78 ± .06

In what degree a limited amount of supervision helps the learning process?

Datasets	Method	0% Anomalies	1% Anomalies	3% Anomalies
KDDCUP99_{Rev}	ARN^N	.99 ± .00	.99 ± .00	.99 ± .00
	ARN^G	.97 ± .01	.99 ± .00	.99 ± .00
	WALDO	-	.80 ± .01	.80 ± .01
Bank	ARN^N	.76 ± .04	.79 ± .06	.89 ± .03
	ARN^G	.70 ± .05	.72 ± .04	.82 ± .05
	WALDO	-	.56 ± .01	.56 ± .01

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Evolution of the outliers

- Can outliers shift into a normal pattern over time?

Evolution of the outliers

- Can outliers shift into a normal pattern over time?

Dynamic Graph Generation for Anomaly Detection
Methodology

- Define a generative model for modeling the evolution of normality patterns
  - Any deviations from the expected behavior are symptoms of outliers
- Graphs are data structures that fit many real scenarios
- FuDGE: a probabilistic generative model for predicting graph evolution.

Dynamic Graph Generation for Anomaly Detection
Methodology

- **Problem**
  - Continuous changes in the graph structure need flexible architectures that can dynamically adjust their dimension over times

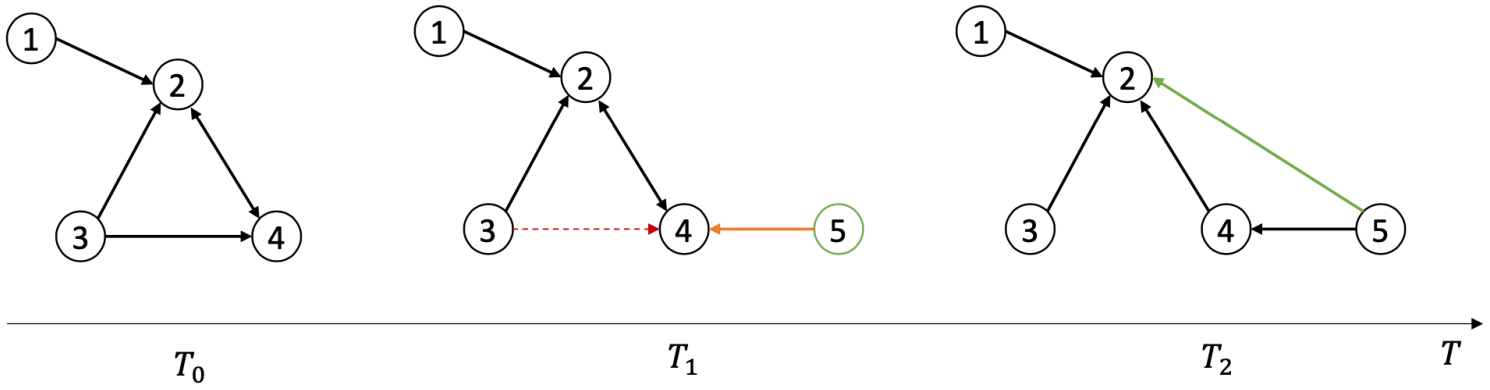

- **Solution**
  - Define a **graph-size invariant** probabilistic generative model for predicting the graph evolution through step-wise changes in the graph structure

Dynamic Graph Generation for Anomaly Detection
Methodology

FuDGE consists of:

- A history network (H) generates a representation of the entire graph history
- A source network and a destination network derive embeddings for both source and destination nodes;
- An auxiliary graph network generates a node embedding used for matching with source/destination embeddings and generating probability scores;
- An event network predicts the event type (addition/deletion).

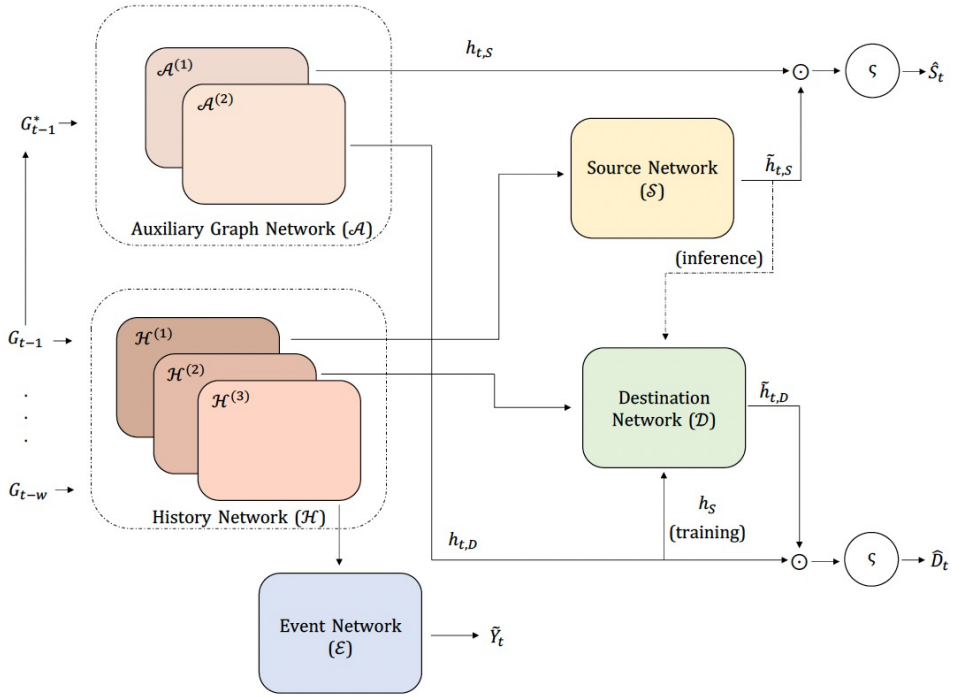

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ1**. Can FuDGE be used to track and predict graph evolution in real-world scenarios? How does it compare to state-of-the-art approaches?

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ1**. Can FuDGE be used to track and predict graph evolution in real-world scenarios? How does it compare to state-of-the-art approaches?

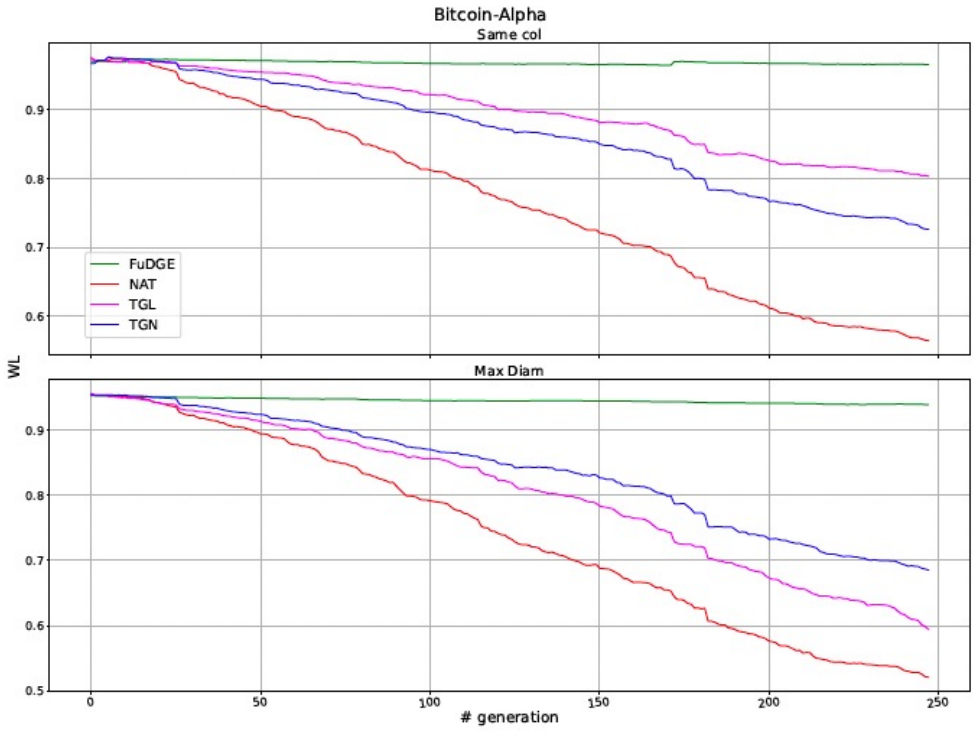

Long-term graph generation

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ2**. How does FuDGE fit real graph features?

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ2**. How does FuDGE fit real graph features?

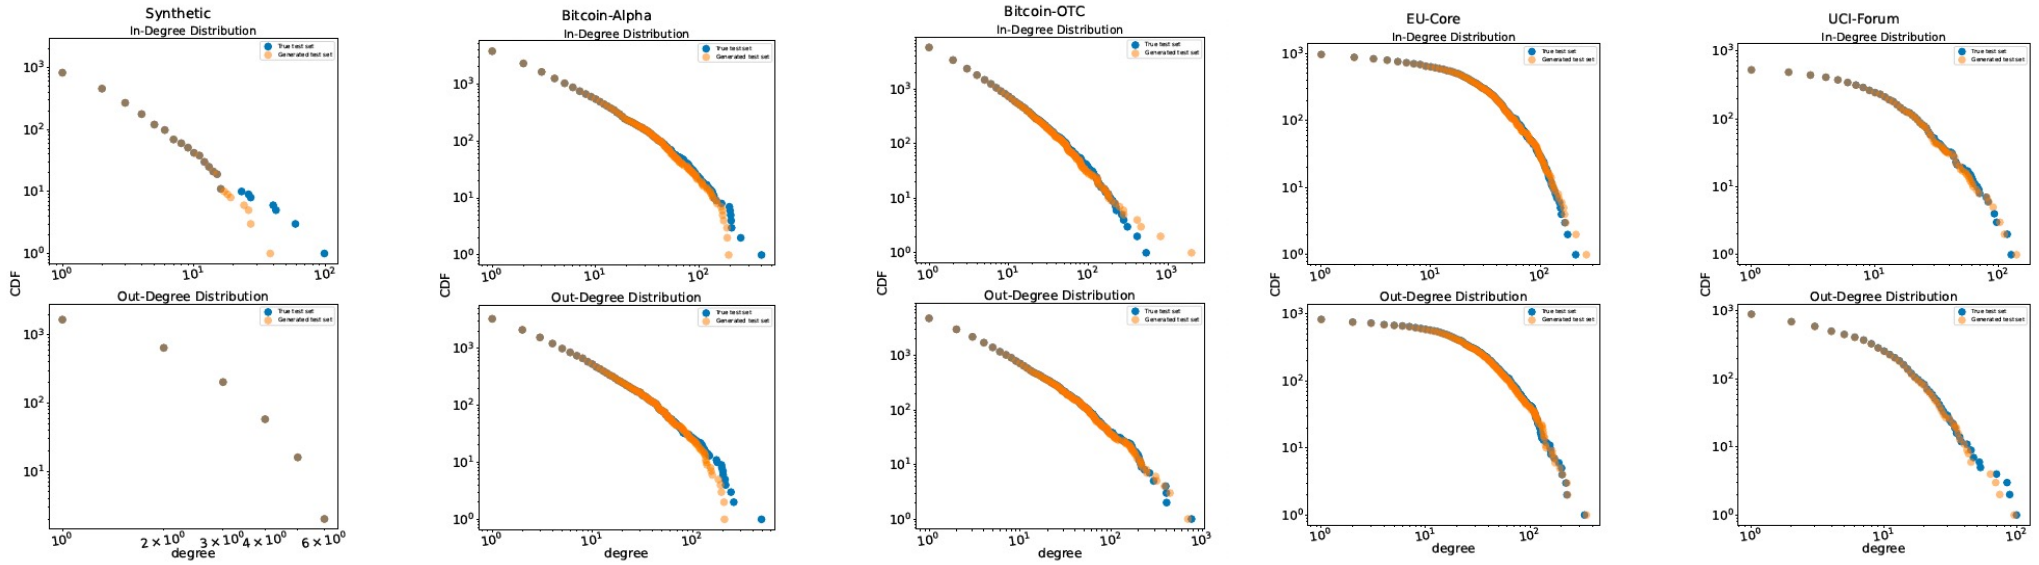

Conclusions

Anomaly Detection & Generation

A Federated-based Anomaly Detection Systems

An Autoencoder-based Anomaly Detection Systems

A Siamese Network for Anomaly Detection

Dynamic Graph Generation for Anomaly Detection

Adversarial Reconstruction Network

About my Academic Career
Publications

1. Comito C., Guarascio M., **Liguori A.**, Pisani F.S., (2024), "Leveraging a Self-Supervised Deep Learning Approach for Detecting Fake News Across Various Domains", Accepted for publication in WSDM – Integrity Workshop
2. Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9
3. Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
4. Comito C., Pisani F.S., Coppolillo E., **Liguori A.**, Guarascio M., Manco G., (2023), "Towards Self-Supervised Cross-Domain Fake News Detection", Italian Conference on Cybersecurity (ITASEC), CEUR Workshop Proceedings, URL: https://ceur-ws.org/Vol-3488/paper12.pdf
5. Coppolillo E., Gallo D., **Liguori A.**, Mungari S., Ritacco E., Manco G., (2023), "Siamese Network for Fake Item Detection", 31st Symposium on Advanced Database System (SEBD), CEUR Workshop Proceedings, vol. 3478, URL: https://ceur-ws.org/Vol-3478/paper57.pdf
6. Coppolillo E., **Liguori A.**, Guarascio M., Pisani F. S., Manco G., (2023) "Generative Methods for Out-of-distribution Prediction and Applications for Threat Detection and Analysis: A Short Review", Digital Sovereignty in Cyber Security: New Challenges in Future Vision. Communications in Computer and Information Science, vol. 1807, DOI: https://doi.org/10.1007/978-3-031-36096-1_5
7. **Liguori A.**, Mungari S., Zuppelli M., Comito C., Cambiaso E., Repetto M., Guarascio M., Caviglione L., Manco G., (2023), "Using AI to face covert attacks in IoT and softwarized scenarios: challenges and opportunities" Ital-IA - 3rd National Conference on Artificial Intelligence, CEUR Workshop Proceedings, vol. 3486, URL: https://ceur-ws.org/Vol-3486/37.pdf
8. **Liguori A.**, Caroprese L., Minici M., Veloso B., Spinnato F., Nanni M., Manco G., Gama J., (2023), "Modeling Events and Interactions through Temporal Processes – A Survey", arXiv, URL: https://arxiv.org/abs/2303.06067
9. Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

About my Academic Career
Publications

10. Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

11. **Liguori, A.**, Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

12. **Liguori, A.**, Manco, G., Pisani, F.S., Ritacco, E., (2021), "Adversarial Reconstruction for Anomaly Detection and Generation", IEEE International Conference on Data Mining (ICDM), DOI: https://doi.org/10.1109/ICDM51629.2021.00145

13. Folino F., Guarascio M., **Liguori A.**, Manco G., Pontieri L., Ritacco E., (2020), "Exploiting Temporal Convolution for Activity Prediction in Process Analytics", in ECML PKDD 2020 Workshops, vol. 1323, DOI: https://doi.org/10.1007/978-3-030-65965-3_17

14. Scicchitano F., **Liguori A.**, Guarascio M., Ritacco E., Manco G., (2020), "A Deep Learning Approach for Detecting Security Attacks on Blockchain", in Italian Conference on Cybersecurity (ITASEC), CEUR Workshop Proceedings, URL: https://ceur-ws.org/Vol-2597/paper-19.pdf

15. Scicchitano F., **Liguori A.**, Guarascio M., Ritacco E., Manco G., (2020), "Deep Autoencoder Ensembles for Anomaly Detection on Blockchain", in International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 12117, DOI: https://doi.org/10.1007/978-3-030-59491-6_43

About my Academic Career
Publications Under Review

- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)
- **Liguori A.**, Caroprese L., Minici M., Veloso B., Spinnato F., Nanni M., Manco G., Gama J., "*Modeling Events and Interactions through Temporal Processes – A Survey*", Submitted: ACM Computing Surveys
- **Liguori A.**, Ritacco E., Benvenuto G., Iiritano S., Manco G., Ruffolo M., "Siamese Networks for Unsupervised Failure Detection in Smart Indutry ", Submitted: 27th International Symposium on Methodologies for Intelligent Systems – Industry Session
- **Liguori A.**, Zuppelli M., Gallo D., Guarascio M., Caviglione L., "Erasing the Shadow: Sanitization of Images with Malicious Payloads using Deep Autoencoders", Submitted: 27th International Symposium on Methodologies for Intelligent Systems
- Guarascio M., **Liguori A.**, Manco G., Ritacco E., "*Knowledge Discovery in Databases*", Submitted: Encyclopedia of Bioinformatics and Computational Biology, 2nd Edition

About my Academic Career
Experiences

- Doctoral Abroad Period at Boise State University (BSU) (Computer Science Department), Boise, ID, United States

- Assistant Professor for Master's Degree Courses (DeMaCS, UNICAL)

Projects

SERICS
SECURITY AND RIGHTS IN THE CYBERSPACE

HUMANE AI

SECURE.
open nets

True Detective 4.0

ICAR
CNR
Istituto di Calcolo
e Reti ad Alte Prestazioni

ICAR-CNR Researcher Associate

Thank You for your attention!

Appendix

Contributions

Anomaly Detection

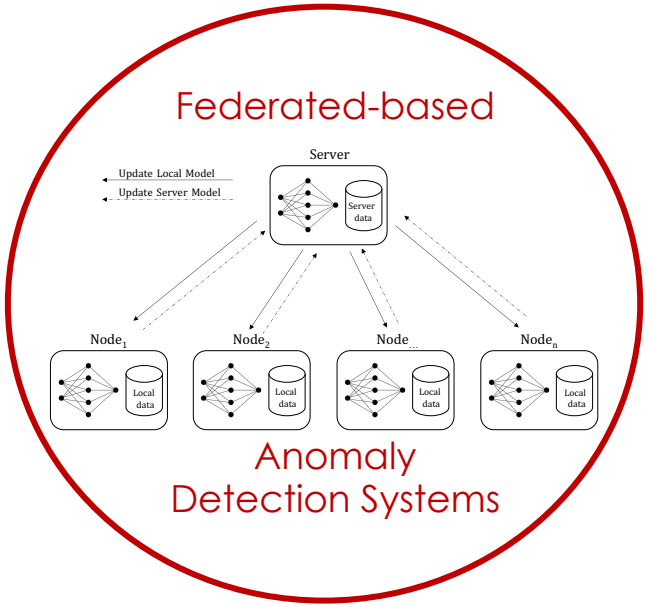

Federated-based Anomaly Detection Systems
Methodology

RQ: Can a Machine/Deep learning-based anomaly detection model be trained without moving data toward a single/centralized device?

A: Yes, Federated Learning is the answer!

Federated learning allows cooperation among different clients/devices to identify anomalous behaviors in data distributed across these entities.

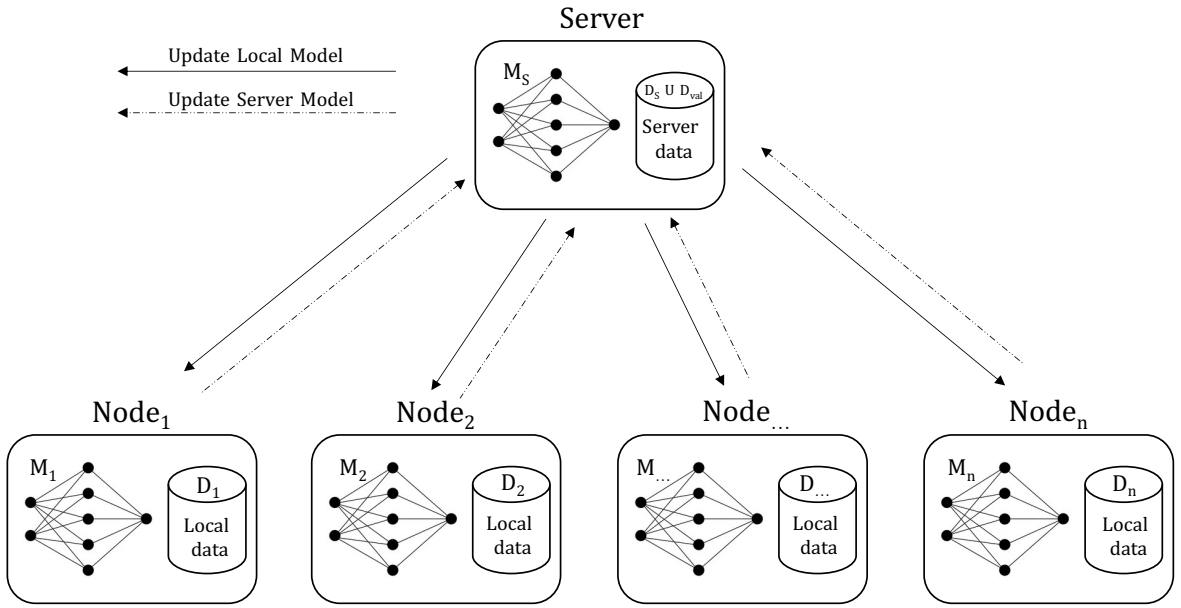

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

Federated-based Anomaly Detection Systems
Experiments

RQ1: How do the performances of the end nodes improve during the federated learning iterations?

RQ2: How is the federated approach effective in detecting anomalous behaviors?

Iteration	Model	AUC	AUPRC	F1-Score
Initialization	*server*	0.926	0.745	0.699
1	$peer_{avg}$	0.943	0.775	0.737
	server	0.926	0.745	0.699
2	$peer_{avg}$	0.955	0.805	0.770
	server	**0.959**	**0.819**	**0.763**
3	$peer_{avg}$	0.955	0.804	0.768
	server	0.959	0.819	0.763
4	$peer_{avg}$	0.960	0.816	0.779
	server	0.959	0.819	0.763
5	$peer_{avg}$	0.960	0.816	0.782
	server	0.959	0.819	0.763
6	$peer_{avg}$	0.959	0.813	0.774
	server	0.959	0.819	0.763
7	$peer_{avg}$	0.960	0.820	0.783
	server	**0.962**	**0.826**	0.641
8	$peer_{avg}$	0.965	0.829	0.797
	server	**0.971**	**0.845**	0.744
9	$peer_{avg}$	0.959	0.818	0.783
	server	0.971	0.845	0.744
10	$peer_{avg}$	0.965	0.829	0.811
	server	0.970	0.842	**0.817**

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

Federated-based Anomaly Detection Systems
Experiments

RQ3: How are the performances of the federated approach compared with a centralized one?

Approach	Coding	AUC	AUPRC	F1-Score
Centralized	Plain	**0.972**	**0.851**	**0.835**
	Base64	**0.899**	**0.605**	0.589
	zip	0.776	0.397	0.344
Federated	Plain	0.970	0.842	0.817
	Base64	0.893	0.594	**0.614**
	zip	**0.856**	**0.498**	**0.363**

FL approach achieves comparable performances with a fully centralized method w/o the necessity of moving data toward a single node.

- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Manco G., Zuppelli M., (2023), "A federated approach for detecting data hidden in icons of mobile applications delivered via web and multiple stores", Social Network Analysis and Mining (SNAM), vol. 13, DOI: https://doi.org/10.1007/s13278-023-01121-9
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Surace G., Zuppelli M., (2022), "Federated Learning for the Efficient Detection of Steganographic Threats Hidden in Image Icons", EAI International Conference on Pervasive knowledge and collective intelligence on Web and Social Media (PerSoM), vol. 494, DOI: https://doi.org/10.1007/978-3-031-31469-8_6

Appendix

Anomaly Detection

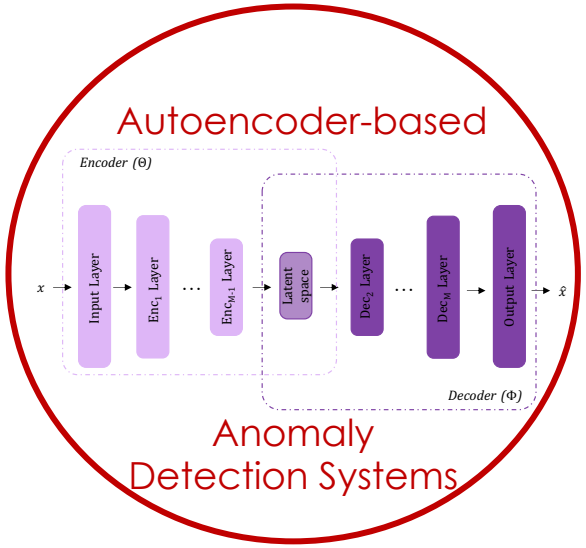

Autoencoder-based Anomaly Detection System
Methodology

RQ: Can we build a model for learning to model normal behavior?

A: Use encoder-decoder architecture for learning to reproduce normal data. Reconstruction error is adopted for detecting anomalies.

Normal data exhibit low reconstruction error, causing anomalies to be elements with a high reconstruction error.

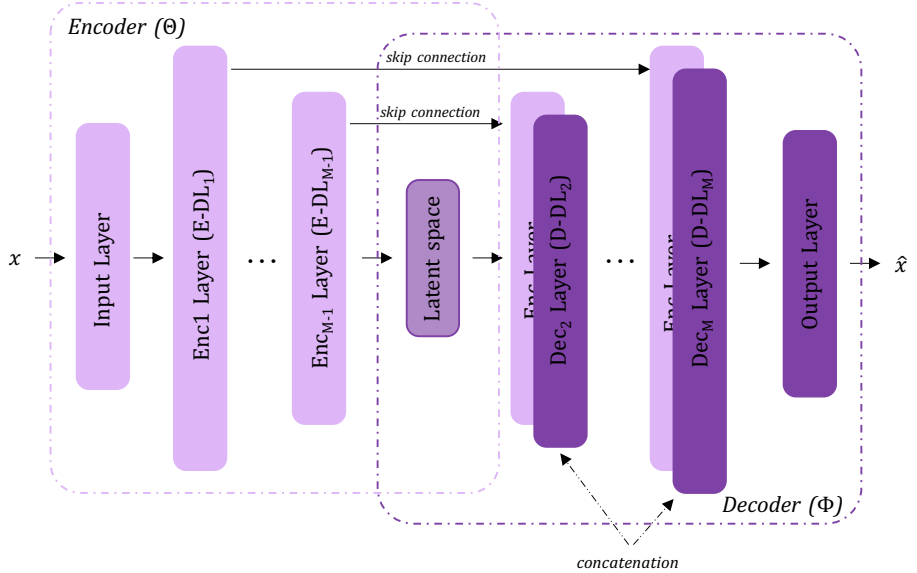

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Autoencoder-based Anomaly Detection System
Methodology

- An incremental learning scheme based on an ensemble of encoder-decoder architectures is implemented
  - Each model is trained using the weights of the previous model
  - The anomaly score is computed by merging the reconstruction errors provided by each base model

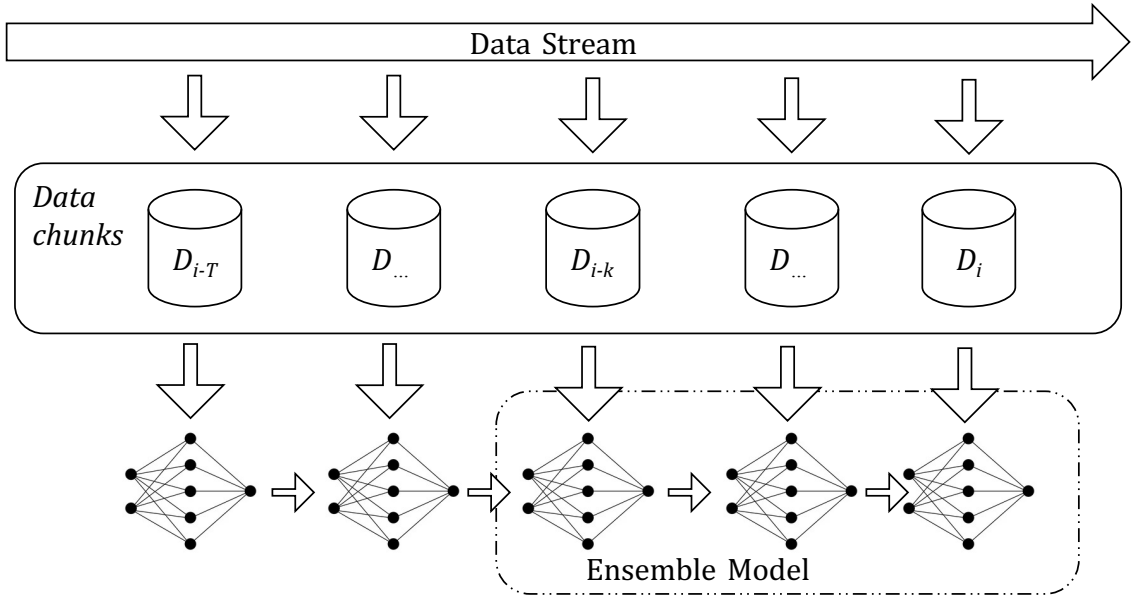

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Autoencoder-based Anomaly Detection System
Experiments

RQ1: How do the different combination schemes and ensemble sizes influence the predictive performances?

Ensemble Size	Strategy	Detection Threshold	Accuracy	Precision	Recall	F1-Score
3	*median*	90^{th} perc.	0.894	0.771	0.979	0.863
		95^{th} perc.	0.947	0.902	0.948	0.924
		99^{th} perc.	**0.955**	**0.950**	0.915	**0.932**
	max	90^{th} perc.	0.894	0.772	0.974	0.861
		95^{th} perc.	0.946	0.901	0.945	0.922
		99^{th} perc.	0.948	0.941	0.902	0.921
	avg	90^{th} perc.	0.892	0.767	0.980	0.860
		95^{th} perc.	0.947	0.901	0.947	0.924
		99^{th} perc.	0.952	0.946	0.910	0.928
5	*median*	90^{th} perc.	0.890	0.764	0.977	0.858
		95^{th} perc.	0.933	0.863	0.954	0.906
		99^{th} perc.	0.952	0.944	0.911	0.927
	max	90^{th} perc.	0.893	0.767	**0.981**	0.861
		95^{th} perc.	0.939	0.878	0.953	0.914
		99^{th} perc.	0.941	0.944	0.878	0.910
	avg	90^{th} perc.	0.891	0.765	**0.981**	0.859
		95^{th} perc.	0.941	0.884	0.951	0.916
		99^{th} perc.	0.951	0.942	0.912	0.927

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Autoencoder-based Anomaly Detection System
Experiments

RQ2: How does the deep ensemble model compare with base deep models? Can the ensemble strategy improve the performance of the 'single' model?

Model Type	Detection Threshold	Accuracy	Precision	Recall	F1-Score
Sparse U-Net	90^{th} perc.	0.882	0.743	**0.993**	0.850
	95^{th} perc.	0.921	0.822	0.976	0.893
	99^{th} perc.	0.936	0.942	0.865	0.902
DAE	90^{th} perc.	0.869	0.724	**0.993**	0.837
	95^{th} perc.	0.910	0.801	0.975	0.880
	99^{th} perc.	0.905	**0.962**	0.750	0.843
Sparse-AE	90^{th} perc.	0.875	0.737	0.979	0.841
	95^{th} perc.	0.901	0.795	0.951	0.866
	99^{th} perc.	0.902	0.922	0.778	0.844
Skip-AE	90^{th} perc.	0.875	0.736	0.982	0.841
	95^{th} perc.	0.907	0.799	0.968	0.876
	99^{th} perc.	0.819	0.853	0.563	0.678
Ensemble (k=3)	90^{th} perc.	0.894	0.771	0.979	0.863
	95^{th} perc.	0.947	0.902	0.948	0.924
	99^{th} perc.	**0.955**	0.950	0.915	**0.932**

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Autoencoder-based Anomaly Detection System
Experiments

RQ3: How is the proposed model robust to lack of data, i.e., limited training data?

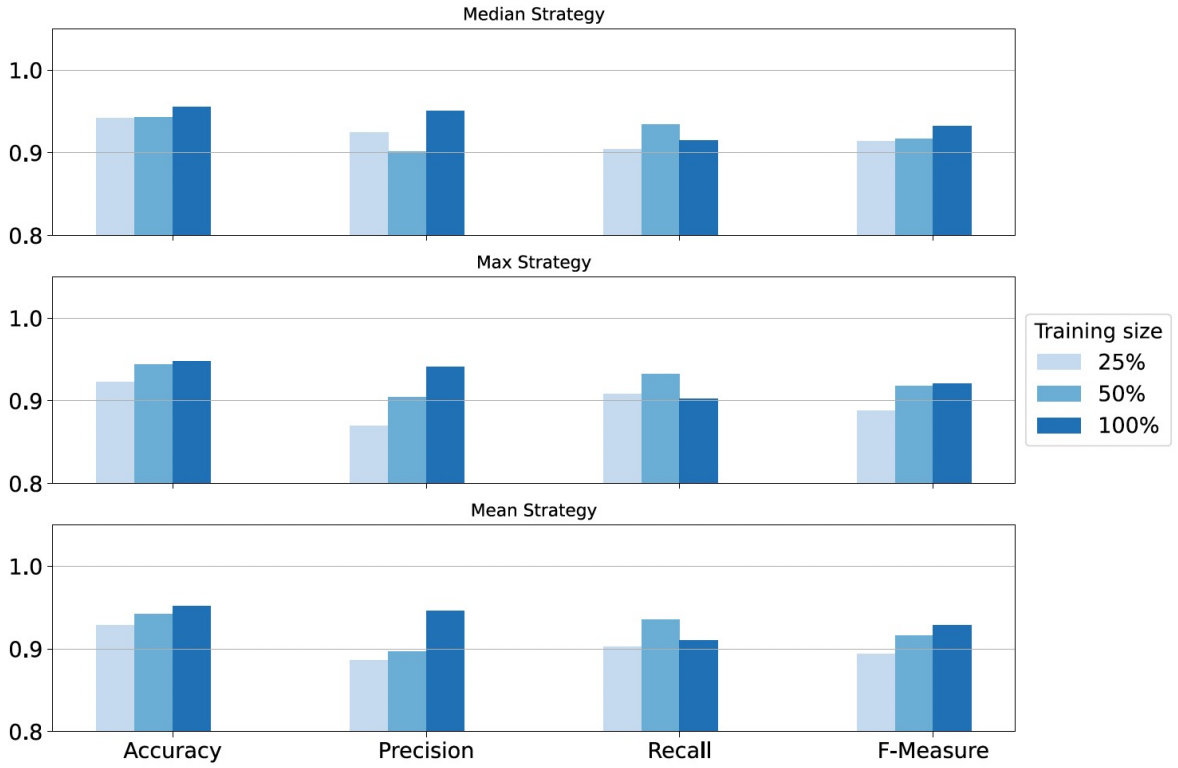

- Cassavia N., Caviglione L., Gurascio M., **Liguori A.**, Zuppelli M., (2023) "Learning Autoencoder Ensembles for Detecting Malware Hidden Communications in IoT Ecosystems", Journal of Intelligent Information Systems (JIIS), DOI: https://doi.org/10.1007/s10844-023-00819-8
- Cassavia N., Caviglione L., Guarascio M., **Liguori A.**, Zuppelli M., (2022) "Ensembling Sparse Autoencoders for Network Covert Channel Detection in IoT Ecosystems", International Symposium on Methodologies for Intelligent Systems (ISMIS), vol. 13515, DOI: https://doi.org/10.1007/978-3-031-16564-1_20

Appendix

Anomaly Detection

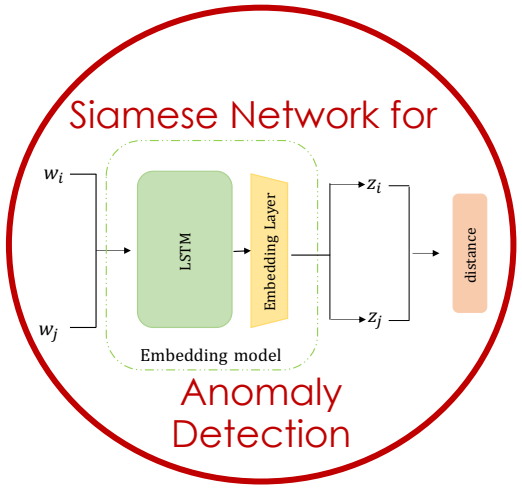

Siamese Neural Network for Anomaly Detection
Methodology

RQ: Can we define a model that can learn different normality behavior modes for isolating outliers?

A: Use a Siamese Network for mapping data into data points lying on a latent space

- Data sharing the same/similar characteristics are close in the latent space, meaning they belong to the same category, i.e., a possible normality pattern.
- All the data placed far from the different normality patterns are considered anomalies.

Liguori, A., Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

Siamese Neural Network for Anomaly Detection
Methodology

- Apply a sliding window procedure for generating a set of fixed-size observation windows
- Make the Siamese network compare pairs of subsequences: The network learns how to distinguish sub-sequences belonging to the same micro-category

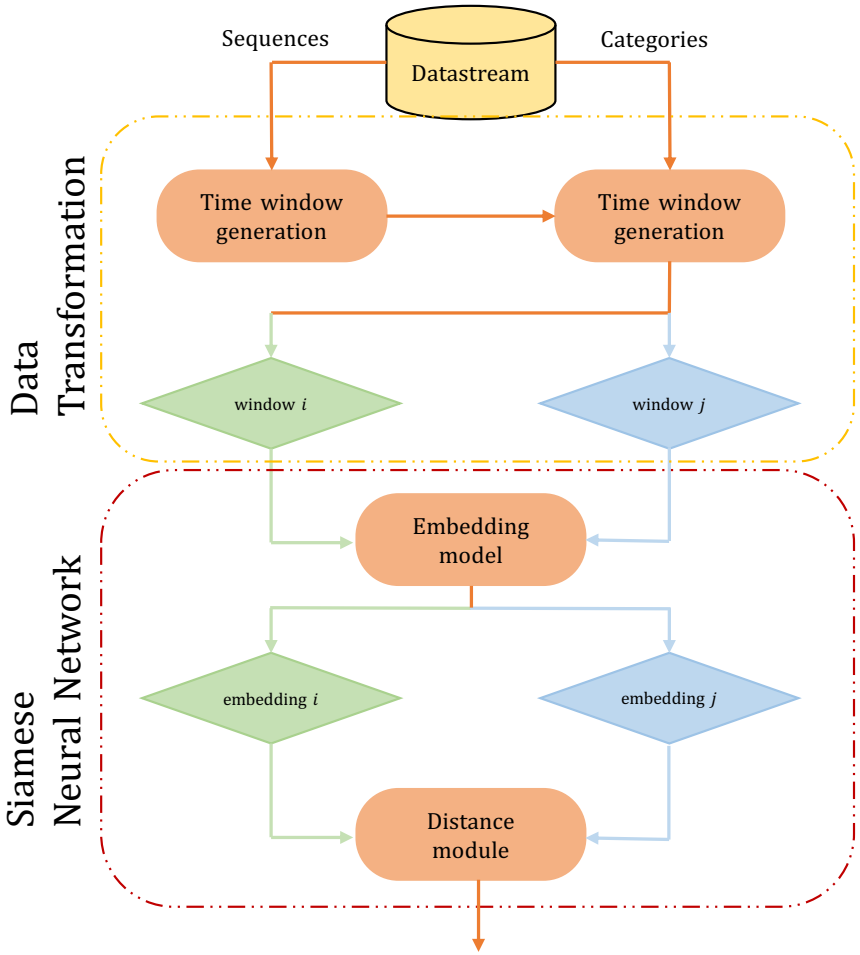

Liguori, A., Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

Siamese Neural Network for Anomaly Detection
Methodology

- Since anomalies are rare events, the network will configure its parameters mainly according to regular data
- When the network finds something that is distant from all the discovered normality patterns, it is classified as anomalous.

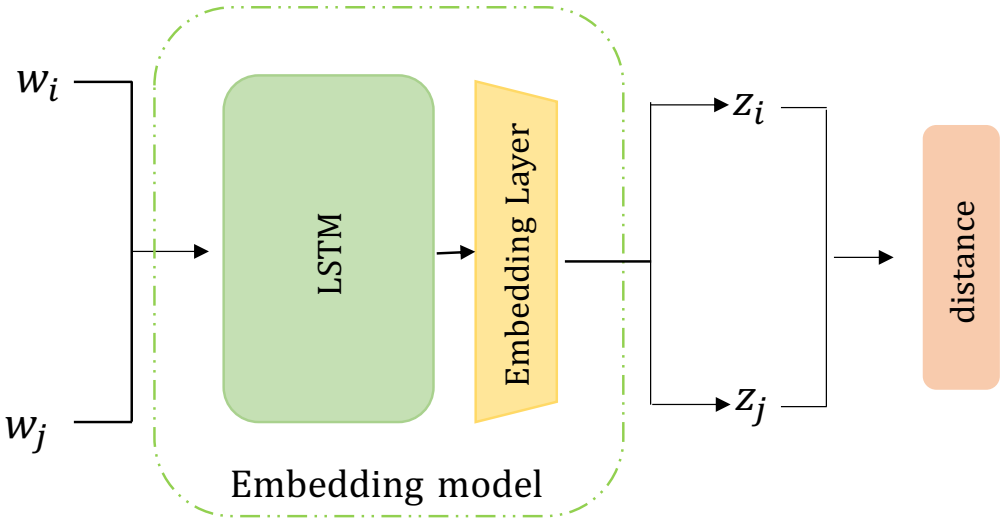

Liguori, A., Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

Siamese Neural Network for Anomaly Detection
Experiments

RQ1. Is the model able to perfectly separate the different normality patterns?

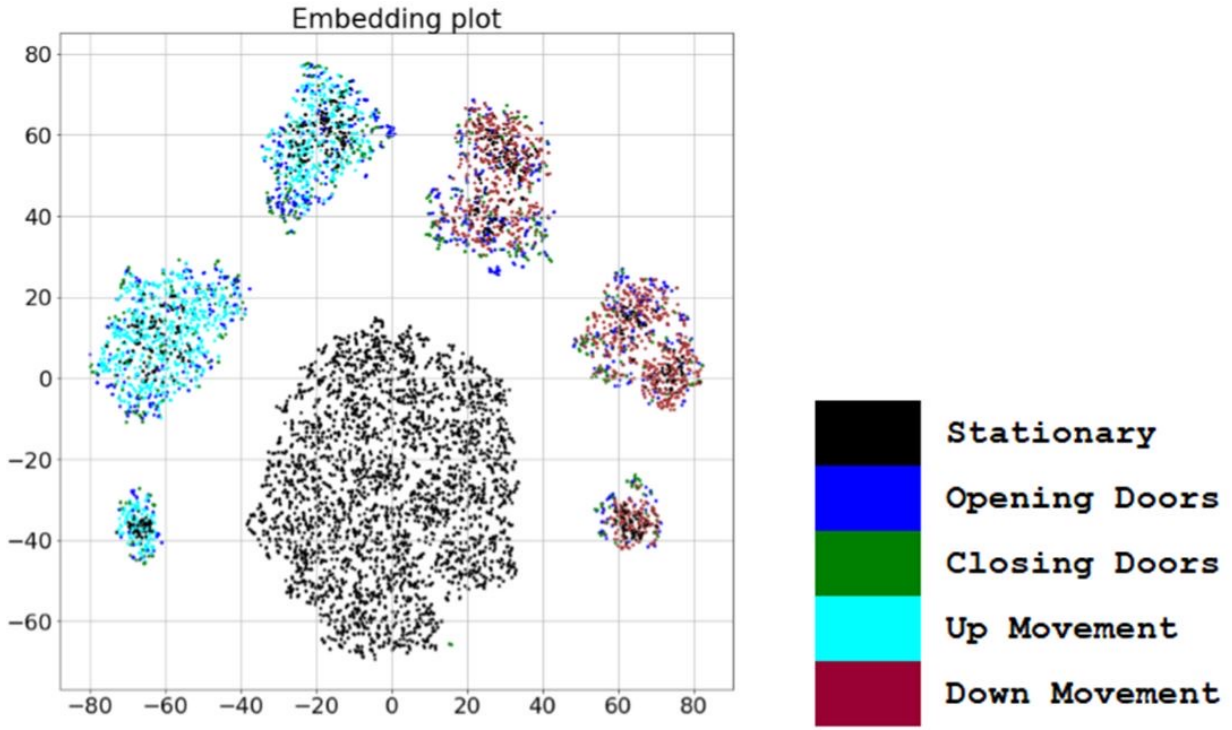

Liguori, A., Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

Siamese Neural Network for Anomaly Detection
Experiments

RQ2. Is the Siamese-based approach effective in detecting anomalous situations?

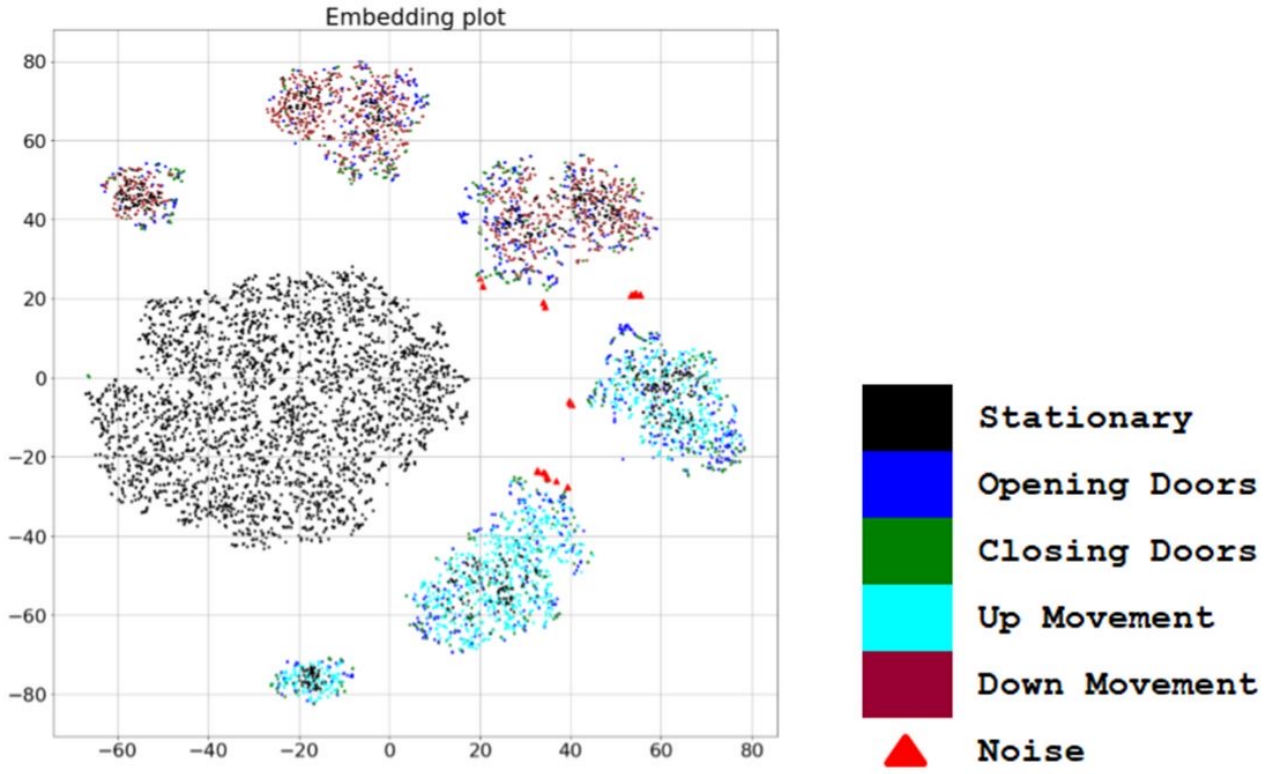

Liguori, A., Manco, G., Ritacco, E., Ruffolo, M., Iiritano., S., (2021), "A Deep Learning Approach for Unsupervised Failure Detection in Smart Industry (Discussion Paper)", The 29th Italian Symposium on Advanced Database Systems, SEBD 2021, URL: https://ceur-ws.org/Vol-2994/paper54.pdf

Appendix

Anomaly Generation & Detection

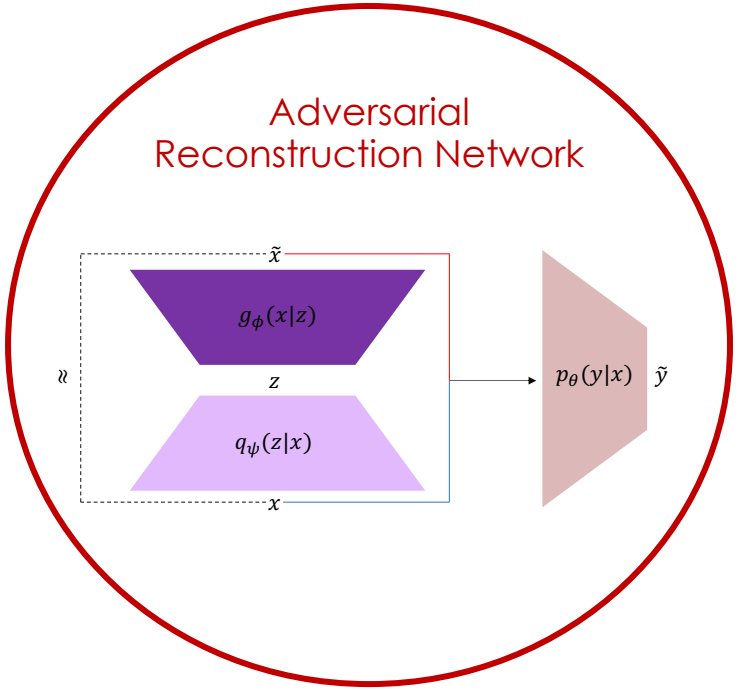

Adversarial Reconstruction Networks
Methodology

The adversarial game has associated the discriminator loss

$$\mathcal{L}_{\mathcal{D}}(\theta|\phi,\psi) = \mathbb{E}_{x\sim\mathbb{P}_{\mathcal{D}}}[\log p_\theta(0|x)] + \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p_\theta(1|\tilde{x})]$$

and the generator loss

$$\mathcal{L}_G(\phi,\psi|\theta) = \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p_\theta(0|\tilde{x})] + \mathbb{E}_{\substack{x\sim\mathbb{P}_{\mathcal{D}} \\ z\sim q_\psi(\cdot|x) \\ \tilde{x}\sim g_{\phi(z)}}}[\log p(x|\tilde{x})] - \mathbb{KL}\big[q_\psi(z|x)||p(z)\big]$$

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Adversarial Reconstruction Networks
Experiments

RQ4. Which components of the model contribute to the overall quality? How do the architectural choices affect the accuracy of the resulting predictions?

Dataset	ARN^{G}	ARN^{N}	$\text{ARN}^{G-\mathbb{KLD}}$	$\text{ARN}^{N-\mathbb{KLD}}$	ARN^{GE}
KDDCUP99	.99 ± .00	1.00 ± .00	.98 ± .02	.98 ± .02	.74 ± 09
KDDCUP99$_{Rev}$	.97 ± .01	.99 ± .00	.98 ± .01	.96 ± .00	.83 ± .05
KDDCUP99$_{Inv}$	1.00 ± .00	1.00 ± .00	1.00 ± .00	1.00 ± .00	.88± .04
NSL-KDD	.99 ± .00	.99 ± .00	.99 ± .01	.98 ± .00	.74 ± .07
DoH	.99 ± .01	.99 ± .01	.73 ± .01	.79 ± .02	.99 ± .01
DoH$_{Inv}$	.98 ± .01	100. ± .00	.99 ± .00	.99 ± .00	.83 ± .04
CoverType	.94 ± .01	.92 ± .04	.94 ± .01	.94 ± .01	.77 ± .08
CreditCard	-	.99 ± .01	-	.96 ± .05	.75 ± .06
Bank	.77 ± .06	.69 ± .07	.74 ± .07	.63 ± .07	.62 ± .04

- **Liguori A.**, Manco G., Pisani F. S., Ritacco E., "*Adversarial Regularized Reconstruction for Anomaly Detection and Generation*", 2021 IEEE International Conference on Data Mining (ICDM), Auckland, New Zealand, 2021, pp. 1204-1209, doi: 10.1109/ICDM51629.2021.00145.
- **Liguori A.**, Ritacco E., Pisani F.S., Manco G., "*Robust Anomaly Detection via Adversarial Counterfactual Generation*", Submitted: Knowledge and Information Systems (Major review)

Appendix

Anomaly Generation & Detection

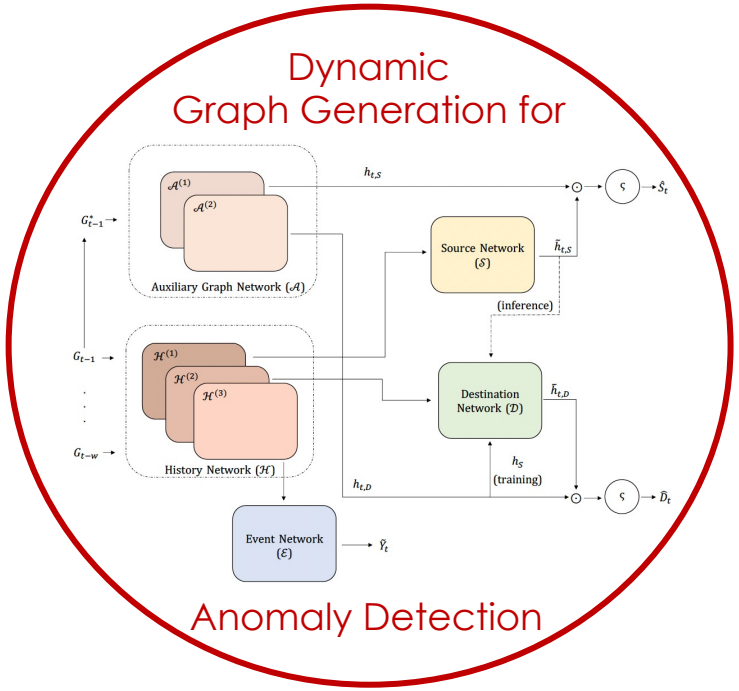

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ1**. Can FuDGE be used to track and predict graph evolution in real-world scenarios? How does it compare to state-of-the-art approaches?

Dataset	Strategy	FuDGE		BA	ER	POWER	TGN			NAT			TGL		
		Configuration	*WL*				*module*	*sample_size*	*WL*	*dimensions*	*sample_size*	*WL*	*variant*	*sample_size*	*WL*
Synthetic	Same col	(i)	**0.98 ± .00**	0.96 ± .00	0.00 ± 0.00	0.97 ± .00	time	4096	0.97 ± .00	[2, 16, 16]	4096	0.97 ± 0.01	JODIE	4096	0.92 ± .00
	Max Diam		**0.94 ± .00**	0.92 ± .00	0.00 ± 0.00	0.93 ± .00	attention	512	0.90 ± 0.02			0.92 ± 0.01	TGAT	512	0.85 ± .02
Bitcoin-Alpha	Same col	(iii)	**0.97 ± .00**	0.50 ± .01	0.01 ± .00	0.55 ± .00	attention	64	0.72 ± .06	[4, 16, 16]	4096	0.69 ± .06	DySAT	512	0.78 ± .03
	Max Diam		**0.94 ± .00**	0.48 ± .01	0.01 ± .00	0.52 ± .01			0.68 ± .06	[2, 32, 16]	64	0.68 ± .08	TGAT	4096	0.59 ± .01
Bitcoin-OTC	Same col	(iii)	**0.92 ± .01**	0.61 ± .00	0.01 ± .00	0.62 ± .00	time	8	0.85 ± .04	[2, 16, 16]	512	0.84 ± .07	DySAT	4096	0.87 ± .02
	Max Diam		0.90 ± .01	0.60 ± .00	0.01 ± .00	0.58 ± .00			**0.87 ± .05**			**0.82 ± .10**	TGN	64	0.78 ± .08
UCI-Forum	Same col	(i)	0.82 ± .01	0.73 ± .00	0.60 ± .01	0.77 ± .00	identity	4096	**0.81 ± .05**	[2, 16, 16]	4096	0.71 ± .04	TGN	4096	**0.86 ± .02**
	Max Diam		0.67 ± .00	**0.67 ± .00**	0.53 ± .00	**0.72 ± .00**		512	0.61 ± .05			0.53 ± .05			**0.72 ± .01**
EU-Core	Same col	(ii)	**0.81 ± .01**	0.48 ± .00	0.15 ± .01	0.45 ± .02	attention	4096	0.53 ± .06	[2, 16, 16]	64	0.43 ± .03	TGAT	512	0.68 ± .09
	Max Diam		**0.76 ± .02**	0.39 ± .00	0.10 ± .01	0.40 ± .00	sum		0.37 ± .06			0.31 ± .04	TGN		0.50 ± .11

Dynamic Graph Generation for Anomaly Detection
Experiments

- **RQ2**. Which components of the model contribute to the overall quality? How do the architectural and learning choices affect the accuracy of the model performance?

Model	Strategy	Configurations		
		(i)	(ii)	(iii)
FuDGE	Same Col	0.9779	0.9757	0.9753
	Max Diam	0.9385	0.9344	0.9348
FuDGE$_{(\mathcal{H},2,\mathcal{A})}$	Same Col	0.9779	0.9642	0.9742
	Max Diam	0.9340	0.9264	0.9328

Table 2: Ablation Study

Dataset	Strategy	70%			60%			50%		
		(i)	(ii)	(iii)	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Synthetic	Same col	**0.9797**	0.9713	0.9592	**0.9666**	0.9461	0.9421	**0.9407**	0.9109	0.9286
	Max Diam	**0.9466**	0.9390	0.9307	**0.9327**	0.9120	0.9000	0.8780	0.8756	**0.8891**
Bitcoin-Alpha	Same col	0.6370	0.8389	**0.9606**	0.5272	**0.9728**	0.9653	**0.5968**	0.5747	0.5935
	Max Diam	0.6666	0.8595	**0.9361**	0.4933	**0.9451**	0.9397	0.4268	0.4368	**0.5946**
Bitcoin-OTC	Same col	0.8566	0.8710	**0.8946**	0.6788	0.6517	**0.9042**	0.6270	0.6876	**0.9027**
	Max Diam	0.8023	0.7855	**0.8479**	0.7980	0.7779	**0.8931**	0.7482	0.7752	**0.8978**
UCI-Forum	Same col	**0.8187**	0.7477	0.7333	0.6848	0.5655	**0.7109**	0.4574	0.7055	**0.7312**
	Max Diam	**0.6553**	0.6007	0.5889	0.5419	0.2930	**0.5623**	0.2158	0.5326	**0.5485**
EU-Core	Same col	0.7305	**0.7312**	0.6427	0.7491	**0.7576**	0.6131	**0.8605**	0.7203	0.6111
	Max Diam	0.6573	**0.6575**	0.5461	0.7971	**0.8035**	0.4923	**0.8310**	0.6291	0.4896

Table 3: Robustness to decreasing amounts of training data. The results show the WL score computed between the real and generated last graph.