

Batch Distillation Data for Developing Machine Learning Anomaly Detection Methods

Justus Arweiler,[†] Indra Jungjohann,[†] Aparna Muraleedharan,[‡] Heike Leitte,[¶]
Jakob Burger,[‡] Kerstin Münnemann,[†] Fabian Jirasek,^{*,†} and Hans Hasse[†]

[†]*Laboratory of Engineering Thermodynamics, RPTU Kaiserslautern,
Erwin-Schrödinger-Straße 44, 67663 Kaiserslautern, Germany*

[‡]*Technical University of Munich, Campus Straubing for Biotechnology and Sustainability,
Laboratory for Chemical Process Engineering, Uferstraße 53, 94315 Straubing, Germany*

[¶]*Department of Computer Science, RPTU Kaiserslautern, Erwin-Schrödinger-Straße 44,
67663 Kaiserslautern, Germany*

E-mail: fabian.jirasek@rptu.de

Abstract

Machine learning (ML) holds great potential to advance anomaly detection (AD) in chemical processes. However, the development of ML-based methods is hindered by the lack of openly available experimental data. To address this gap, we have set up a laboratory-scale batch distillation plant and operated it to generate an extensive experimental database, covering fault-free experiments and experiments in which anomalies were intentionally induced, for training advanced ML-based AD methods. In total, 119 experiments were conducted across a wide range of operating conditions and mixtures. Most experiments containing anomalies were paired with a corresponding fault-free

one. The database that we provide here includes time-series data from numerous sensors and actuators, along with estimates of measurement uncertainty. In addition, unconventional data sources – such as concentration profiles obtained via online benchtop NMR spectroscopy and video and audio recordings – are provided. Extensive metadata and expert annotations of all experiments are included. The anomaly annotations are based on an ontology developed in this work. The data are organized in a structured database and made freely available via doi.org/10.5281/zenodo.17395544. This new database paves the way for the development of advanced ML-based AD methods. As it includes information on the causes of anomalies, it further enables the development of interpretable and explainable ML approaches, as well as methods for anomaly mitigation.

1 Introduction

In chemical process engineering, early and reliable anomaly detection (AD) is highly relevant for preventing costly plant failures and malfunctions¹. Undetected anomalies in chemical processes can pose serious hazards, threatening human health and the environment. They can also cause significant economic losses. To reduce such risks, anomalies must be promptly and reliably identified and mitigated through appropriate anomaly management (AM). The large amount of diverse data generated in chemical processes makes automated AD mandatory. AD methods are therefore implemented in virtually all industrial process control systems. Nevertheless, human operators still play a major role in both AD and AM.

In principle, AD can be considered as a classification task, where, in the simplest case, only two classes are distinguished: normal and abnormal. The design of AD methods is not trivial, as it is characterized by conflicting objectives: good scores in true positive (anomaly detected when it occurs) tend to go hand in hand with bad scores in false positive (anomaly detected in normal operation), and good scores in true negative (no anomaly detected in normal operation) tend to go hand in hand with bad scores in false negative (no anomaly detected even though present). A wide variety of AD methods have been described in the literature. The proposed approaches differ so significantly that even their classification is challenging². A coarse classification distinguishes between methods that rely only on process data (data-driven methods) and others that combine these data with a process model (model-based methods). For an overview of existing classical AD methods, we refer the reader to reviews of the topic, e.g.,¹⁻⁶.

The rapid development of machine learning (ML) has opened up new opportunities in many fields, including chemical engineering, that are being explored by many researchers^{7,7-21}. It is, therefore, astonishing that only comparatively few papers exist in which advanced ML methods, most notably deep learning methods, have been applied to AD in chemical processes²²⁻²⁸. This gap can only partially be explained by the fact that AD on high-

dimensional complex time-series data has only recently attracted wider attention in the ML community, as this field of ML has developed rapidly and yielded a large number of methods that could principally be applied to AD in chemical processes^{28–32}. The main problem that hinders the application of modern ML methods to AD in chemical processes is the lack of suitable data for training, validation, and testing. Although such data are abundant in chemical companies, they are not available to academic researchers. This shortage of accessible process data significantly hinders the development of advanced ML-based AD techniques, which rely heavily on large datasets for training. In contrast, classical AD methods, which primarily use data for testing, are less affected by this limitation².

Up until now, the only larger dataset available for the development, training, and testing of ML-based AD methods for chemical processes has been the Tennessee-Eastman Process (TEP) dataset^{28,33,34}. The TEP represents a fictitious continuous chemical process composed of five key unit operations: a reactor, a product condenser, a vapor-liquid separator, a recycle compressor, and a product stripper. As the TEP data were generated through simulation, they differ significantly from real-world process data. Although some noise was artificially superimposed, the data are still, in essence, deterministic as they are based on a physical process model that reflects only some features of real processes. As a result, the data are considerably better behaved than real process data. We have recently applied advanced ML-based AD methods on the TEP data and have obtained excellent results²⁸. However, preliminary tests using data from the plant described in the present paper, applying the same AD methods, have yielded only very poor results. This stark contrast highlights the critical need for real process data for the development and evaluation of ML-based AD techniques. Furthermore, in the ML literature, deep learning architectures, such as Long Short-Term Memory (LSTM) networks, convolutional neural networks (CNNs), autoencoders, and, more recently, transformer-based models, have been explored for AD in time-series data. While these approaches often show strong performance on synthetic or benchmark datasets, their effectiveness remains largely unvalidated on real industrial process data^{22–28}.

Therefore, we provide such data in the present work. An obviously crucial question is thereby: can suitable data be obtained in an academic environment? We argue that this is possible and may even offer advantages over data from full-scale production sites. Most notably, academic environments provide greater flexibility in designing and executing experimental scenarios, and typically allow for more comprehensive documentation. The most basic decision in planning such experiments is choosing a suitable demonstration process. We chose distillation for three reasons: first, it is a very widely applied separation process in the chemical industry; second, it lends itself to scale-up: we can operate plants on small scales (as we have to in the academic environment) and still learn from them about the behavior of full-scale production plants; and third, there are excellent physical models for distillation, which can be leveraged for data augmentation and the creation of hybrid datasets for ML model training in future work.

Chemical production processes are commonly classified as either batch or continuous. We report here on batch distillation, but this project was carried out in the frame of a Research Unit (see [Acknowledgements](#), in which data from a continuous distillation plant are also generated). The batch distillation plant described here is a small, flexible laboratory plant; however, it can emulate most types of faults that occur in industrial-scale operations³⁵. The primary purpose of the present paper is to describe the laboratory-scale batch distillation plant and its operation in sufficient detail to contextualize the generated process data. We also provide information on the types of anomalies that were generated. Furthermore, we present the first extensive database of real batch distillation process data. This paper establishes the foundation for follow-up publications in which this database will be enlarged and diversified with new process data generated with the described batch distillation plant. We openly share our data with the community, which can be accessed via doi.org/10.5281/zenodo.17395544, to foster progress in ML-based AD in chemical engineering.

2 Batch Distillation Plant

2.1 Overview

The laboratory-scale batch distillation plant used to generate dynamic process data is a modified version of the Iludest LM 2/S glass plant. Figures 1 and 2 show the P&I diagram and a photograph of the plant, respectively. A complete list of the equipment, measurement instrumentation, and controls, including the respective suppliers, is provided in the Supporting Information. It also includes a mapping to the labels in the P&I diagram of Figure 1 which are used in the discussion below. The different types of sensors and control equipment are classified here according to the Sensor, Observation, Sample, and Actuator (SOSA) ontology³⁶.

The setup comprises a stirred 2 L reboiler vessel (V001) with a distillation glass column (inner diameter 30 mm) on top. The column has three sections (C001-C003), each equipped with structured laboratory packing (Sulzer DX30, 37.5 cm height), corresponding to approximately 12 theoretical stages for the entire column. The column segments are thermally insulated by an evacuated mirrored double jacket. Heat is supplied to the reboiler vessel V001 by a heating jacket (H701, 350 W) and a submerged electric heating element (H002, 400 W). The column and the reboiler vessel are jacketed, with heating elements (H702-H708) wrapped in thermal insulation to prevent heat loss through the jackets. The overhead vapor is condensed and subcooled in glass heat exchangers (HE001, HE002) and then collected in a buffer vessel (V002). Cooling water is supplied by a thermostat (TCU1) at 283.15 K. Two gear-type pumps (P701, P702) deliver the condensate either to the distillate receiver (V003) or, as reflux, back into the top of the column. Before the subcooled reflux is returned to the column, it is warmed up in the condenser (HE003). The pressure in the plant is subambient and controlled using a membrane vacuum pump (P301). To protect the vacuum pump from solvent residues, a condensation trap (HE004) is used, operated with liquid ethanol

at 233.15 K supplied by a thermostat (TCU2). The plant can be purged with nitrogen. In the experiments of the present work, the plant was operated at fixed pressure, reflux ratio R , and heat supply. The reflux ratio R is the ratio of the mass flow rates of the distillate pumped back into the distillation column, $\dot{m}_{F704}$, and that pumped into the distillate receiver V003, $\dot{m}_{F703}$:

$$R = \frac{\dot{m}_{F704}}{\dot{m}_{F703}} \quad (1)$$

As a constant liquid level in the buffer vessel V002 is maintained, the sum of $\dot{m}_{F703}$ and $\dot{m}_{F704}$ is basically equal to the total mass flow of the condensate (neglecting the small losses into the condensation trap HE004).




Figure 1: P&I diagram of the batch distillation plant. Color code of the lines: product (red), cooling water (dark blue), cooling ethanol (light blue), nitrogen (yellow), and pressure control (green). The equipment list, in which the labels are declared, is given in the Supporting Information.



Figure 2: Photo of the laboratory batch distillation plant. The reboiler vessel (V001) and the distillation column (C001-C003) are surrounded by insulation. The visible glass apparatuses are the condenser (HE001-HE003), the buffer vessel (V002), and the distillate receiver (V003). The pumps (P701, P702) are blue, the condensation trap (HE004) has a yellow jacket and is connected to the gray vacuum pump (P301). The benchtop NMR spectrometer is on the top left rack.

2.2 Actuators

Flexible operation of the batch distillation plant is achieved through the use of diverse actuators, including manually operated components and electronically controlled elements. These actuators serve various purposes, such as controlling fluid flow, heat supply, and pressure. Manually operated actuators, such as sampling valves and ventilation ports, are primarily used for system maintenance and process interventions that require human oversight. In contrast, the actuators integrated into the process control system (PCS, cf. Section 2.4) enable automated, precise, and reproducible operation of the plant. A list of actuators that can be controlled through the PCS is given in Table 1.

Table 1: List of actuators in the batch distillation plant that are integrated in the process control system. The labels are the same as in the P&I diagram in Figure 1.

Equipment type	Label	Controlled variable	Unit	Description
Heater	H002	Power	%	Power supply to the heating rod submerged in reboiler vessel V001
	H701	Power	%	Power supply to the heating jacket of reboiler vessel V001
	H702	Power	%	Power supply to the heating jacket of reboiler vessel V001
	H704	Power	%	Power supply to the heating jacket of column section C001
	H706	Power	%	Power supply to the heating jacket of column section C002
	H708	Power	%	Power supply to the heating jacket of column section C003
Pump	P201	Flow rate	ml min ⁻¹	Volumetric flow rate of liquid to online NMR
	P301	State	on/off	Vacuum pump for system pressure regulation
	P701	Power	%	Power supply of reflux pump
	P702	Power	%	Power supply of distillate withdrawal pump
Valve	AV708	State	open/closed	State of valve at vacuum pump line
	AV716	State	open/closed	State of valve at cooling water line
	TV1	Open cross section	%	Open cross section of the vacuum throttle valve

2.3 Sensors

All conventional sensors of the plant are listed in Table 2. Table 3 gives an overview of the used sensor types and their respective measurement uncertainty. The temperatures of the reboiler vessel, the condenser, and the three column sections, as well as the temperatures of each electrical heating element, are tracked. Moreover, flow rates of the product, reflux, and cooling water, and the levels of the liquid in the reboiler vessel V001 and the buffer vessel V002 are tracked. The pressure in the condenser and the differential pressure between the bottom and the top of the column are monitored. Information on sensor calibration is provided in the Supporting Information.

Table 2: List of sensors of the batch distillation plant. The labels are the same as in the P&I diagram in Figure 1. Sensors marked with a dagger are part of a control loop.

Measured variable	Label	Unit	Feature of interest
Level	LS701	-	Contact to liquid (0: dry, 1: wet) in reboiler vessel V001
	LS702	-	Contact to liquid (0: dry, 1: wet) in buffer vessel V002
Mass flow	FT201	kg h ⁻¹	Liquid sample stream to online NMR
	FT703†	kg h ⁻¹	Top product stream
	FT704†	kg h ⁻¹	Reflux stream
	FYI702	s ⁻¹	Cooling water stream
Pressure difference	PDI701	mbar	Pressure drop along column (reboiler vessel V001 to top vessel G001)
	PDI702†	mbar	Static pressure in buffer vessel V002
Pressure	PY23†	mbar	Pressure in top vessel G001
Temperature	T701	°C	Heating jacket H701 for lower hemisphere of reboiler vessel V001
	T702†	°C	Heating jacket H701 for upper hemisphere of reboiler vessel V001
	T703†	°C	Liquid residue in reboiler vessel V001
	T704†	°C	Heating jacket H704
	T705†	°C	Fluid on top of column section C003
	T706†	°C	Heating jacket H706
	T708†	°C	Heating jacket H708
	T709	°C	Fluid beneath column section C001
	T711	°C	Fluid between column sections C001 and C002
	T712	°C	Fluid between column sections C002 and C003

Table 3: Sensor types and measurement uncertainties. The labels are those of the corresponding components in the P&I diagram in Figure 1. The actual measurement uncertainty may depend on the measured value.

Measured variable	Labels	Sensor type	Uncertainty
Level	LS701, LS702	Diffuse-reflective photoelectric sensor	n.a.
Mass flow	FT703, FT704	Coriolis mass flow sensor	$\pm 0.06\%$
	FT201	Coriolis mass flow sensor	$\pm 0.2\%$
	FYI702	Paddlewheel flow sensor	$\pm 15\%$
Pressure difference	PDI701, PDI702	Piezoresistive sensor	$\pm 1\%$
Pressure	PY23	Capacitance manometer sensor	$\pm 0.05\%$
Temperature	T701, T702, T704,	Pt100 class A	± 0.95 K
	T706, T708		
	T703, T705, T709,	Pt100 class 1/3	± 0.1 K
	T711, T712		

In addition to the conventional sensors described above, the plant is equipped with several unconventional sensors. Three Logitech BRIO Ultra-HD Pro Business webcams are placed at different locations within the plant to capture high-resolution images every 2 seconds. Sample images are shown in Figure 3. One camera (CAM001) monitors the glass condensers HE001 and HE003, including the condensate drain and the reflux inlet. The two other cameras monitor the filling levels of the buffer vessel V002 (CAM002) and the distillate receiver V003 (CAM003). Furthermore, audio signals from plant operations are collected with a RØDE AI-Micro microphone, which was either placed beneath the reboiler vessel V001 or between the reflux and distillate withdrawal pumps P701 and P702.

Online concentration measurements were carried out with a Spinsolve 80 ULTRA Carbon benchtop NMR spectrometer (QR201) from Magritek operated with the Magritek software “Spinsolve” Version 2.3.8. Liquid from either the reboiler vessel V001 or the buffer vessel V002 was continuously pumped through a flow cell in the NMR spectrometer using pump P201. Details on the online NMR measurements and the data analysis are given in the Supporting Information.

In addition to the online NMR analysis, samples were also analyzed offline by gas chromatography. Liquid samples were drawn from the reboiler vessel V001 and the buffer vessel V002 using a gas-tight syringe to prevent air from entering the plant during sampling. The samples were collected approximately every 20 minutes during the experimental runs. Further details on the offline concentration analysis are given in the Supporting Information.




Figure 3: Snapshots from the cameras in the batch distillation plant. Left: Glass heat exchangers HE001 and HE003. Middle: Buffer vessel V002 with connections. Right: Distillate receiver V003.

2.4 Process Control

The batch distillation plant is operated using a Python-based process control system (PCS) implemented in the itom software package³⁷, version 4.2.2, an open-source integrated development environment (IDE) for automation. Communication between the PCS and the batch distillation plant is realized using the OPC/UA data exchange format, with software from Iludest called “IluLab” acting as an intermediate communicator. The PCS directly communicates with the electronic lab journal “elabFTW”³⁸, which is also used for planning the experimental campaigns and storing the metadata for each experiment.

Before a distillation experiment is conducted, the corresponding experimental run is created in the electronic lab journal. All sensor data from the PCS are read and stored every second. Furthermore, all events in the PCS, e.g., the start and stop of experiments, changes in actuator settings, violations of security thresholds, or moving to the next step in an experimental recipe, are stored in a log file with corresponding timestamps. Metadata stored

for each experiment in the electronic lab journal include ambient conditions, as described in the Supporting Information, anonymized information about the operators conducting the experimental run, and information on the plant. Changes to the setup or control system, and the operating times of each plant component, are tracked in the electronic lab journal. Hence, the currently installed version of the plant, the feed mixture, the operating point, and information on anomalies are linked to the run. Furthermore, a distillation recipe containing start-up and shut-down instructions for the PCS and operating point settings is linked to the experiment.

For each experiment, the corresponding distillation recipe is loaded into the PCS, enabling standardized procedures for conducting experiments and for plant operation at defined operating points. The operating point of the plant is fixed by set values for the absolute pressure at the column top (PY23), the heat supply to the reboiler vessel (H701 and H002), and the reflux ratio R . Thereby, the maximum pressure is atmospheric, and the maximum heat supply is 750 W. The control values for the actuators are set based on the operating point.

Several actuators are operated using PI and PID controllers; detailed information about the controllers is given in Table 4.

Table 4: Overview of control circuits used in the batch distillation plant. The labels are those of the components in the P&I diagram in Figure 1.

Actuator	Target	Type
H702	T703	PI
H704	T703	
H706	T703	
H708	T705	
P701	PDI701	PID
P702	PDI702	
TV1	PY23	
P201	FT201	

3 Operation

3.1 General Procedure

Before each experiment, a leakage test was performed: the dry, empty plant was sealed, and the pressure was set to $p_{PY23} = 10$ mbar. After reaching this pressure, the pressure control was stopped, and the pressure was tracked. The test was considered to be passed if the pressure increase was below 0.5 mbar per hour.

Before the start of the experiment, the plant was flushed with nitrogen, and 1-1.5 kg of the liquid feed mixture of known composition were filled into the reboiler vessel V001. Three phases of the experiment are distinguished here: start-up, operation (either normal or with anomalies), and shut-down; an example is shown in Figure 4. During start-up, first, the desired operating pressure is set. When the set pressure is attained, the primary heat sources, H701 and H002, are ramped up to the specified power. The liquid in the reboiler vessel is first heated up before it begins to boil. The rising vapors cause a rapid increase in temperature in the column sections C001-C003. The overhead vapor condenses in HE001 and HE003. The condensate is collected in the buffer vessel V002 and, during start-up, is entirely returned to the column as liquid reflux. The start-up of the process ends once distillate is drawn into the distillate receiver V003. The withdrawal of distillate begins when the temperatures in the distillation column stop shifting significantly over time.

During the operation phase, the condensate in the buffer vessel V002 is split into reflux and product, collected in the distillate receiver V003. Eventually, anomalies are introduced in the operation phase, as described in more detail below. The operation phase ends either after a specified amount of distillate is obtained (or, equivalently, the filling level of the reboiler vessel drops under a certain value) or simply upon time-out. Then, the shut-down phase begins.

For the shut-down, first all heating elements were ramped down, and the pressure control was

switched off. Online sampling continued at least until all samples drawn from the process by the inline-sampling pump P201 during the operation phase were measured by NMR. After all temperatures in the plant had returned to ambient, the ventilation valve HV007 was opened, leading to ambient pressure in the plant. The reboiler vessel V001 was then emptied and rinsed with suitable cleaning agents, such as 2-propanol or purified water. Moreover, the buffer vessel V002 and the distillate receiver V003 were emptied. The NMR sample line was rinsed and emptied. Finally, the plant was sealed and dried under vacuum for 6 to 12 hours using the vacuum pump P301. The plant is considered to be sufficiently clean if it passes the leakage test in the following experiment. The described procedure was applied when the mixture was changed in the following experiment. If the same mixture was studied, a simplified cleaning procedure was applied.




Figure 4: Temperatures and pressure recorded during experiment `batch_dist_ternary_butan-1-ol+propan-2-ol+water/operating_point_001/test_abormal_experiment_001`. The three phases, start-up, operation, and shut-down are indicated.

3.2 Anomalies

Perturbations were intentionally and systematically induced in the operating plant to generate anomalous process data. The perturbations disrupt normal operation, leading to sensor

readings that deviate from expected behavior. Any such observable deviation recorded by the sensors is called an anomaly here. The perturbations were introduced at a specific point in time (typically by changing the setting of some actuators) and then maintained for a specific time span before being removed again (typically by resetting the actuators to their initial values). A variety of perturbations were induced, and their timing and duration were varied. They include well-known industrial failure causes³⁵ and can be categorized as follows: (i) Setpoint changes of actuators, e.g., heaters, thermostats, pumps, or automatic valves, (ii) compromising sensor data, e.g., by noise or drift, and (iii) addition or removal of substances from the plant, e.g., foaming agents or nitrogen. An overview of the induced perturbations is given in Table 5.

Table 5: Overview of the induced perturbations. For the three categories, the target of the perturbation and the manipulated components are reported.

Category	Target	Component
Setpoint changes of actuators	Main heat supply to reboiler vessel (V001)	H002, H701
	Heat supply to compensate heat losses in V001, C001-C003	H702, H706, H704, H708
	Condenser cooling (HE001, HE002)	AV716, TCU1
	Pressure control	TV1
	Reflux ratio	P701, P702
Compromising sensor data	Control loop between sensor and actuator	PDI701, PDI702, T703, T705, T709, T711, T712
Addition or removal of substances	System boundary	HV001-HV005, HV009
	Fluid dynamics	V001, C001-C003

For most experiments with an anomaly, we have also carried out a corresponding fault-free experiment. Furthermore, unintended anomalies occurred in a few experiments. These experiments are also included in the database, but they are not discussed in the present section.

In most cases, the plant recovered after the perturbation was removed. In these cases, the response of the plant to the perturbation can be split into three phases (see Figure 5 for an example):

- Blind phase: The perturbation is induced, but its effects are not yet observable.
- Anomalous phase: This phase begins when the anomaly becomes observable in any of the sensor readings and ends when the perturbation is removed.
- Recovery phase: This phase begins when the perturbation is removed and ends when its effects are no longer observable.

Depending on the experiment, the time the perturbation was applied ranged from 2 to 20 minutes, and the time it took the plant to recover ranged from a few seconds to half an hour.




Figure 5: Temperatures and pressure recorded during experiment `batch_dist_ternary_butanol+propan-2-ol+water/operating_point_003/test_abormal_experiment_002`, in which a perturbation was introduced at $t = 69$ min and removed at $t = 76$ min. The colors indicate the three phases of the response to the perturbation. The perturbation was a setpoint change of the pressure from $p = 500$ mbar to $p = 700$ mbar.

Furthermore, the database includes a few experiments in which the perturbation remained active until the end of the run. An example is shown in Figure 6.




Figure 6: Temperatures recorded during experiment `batch_dist_ternary_butan-1-ol+propan-2-ol+water/operating_point_013/test_abormal_experiment_001` in which a perturbation was introduced at $t = 50$ min. The process did not recover from the perturbation, which was a fault in the cooling water supply that resulted in the dry-running of the reboiler vessel V001. Hence, the process needed to be shut down. The colors indicate the two occurring phases of the response to the perturbation.

4 Data

4.1 Overview of Experimental Data and the Metadata

During each experiment, time-series data are collected. Each experiment generates time-series data over a period of 1-8 h, with a sampling rate of 1 s. In the database, we provide actuator and sensor data of each experiment in separate files. For each sensor, information on the measurement uncertainties is also provided. The time-series data is labeled to distinguish start-up, operation, and shut-down. Missing data points, e.g., due to timing issues in the PCS, were interpolated, resulting in .csv files containing time-series data from 18 sensors and 13 actuators for each second of an experiment. The time-series anomaly labels for the operation are provided in separate .csv files. Structured metadata provide context for the anomalous process data by specifying the occurring anomalies as described in Section 4.3.

Moreover, liquid density, pressure, and mass flux data obtained from the NMR sample pump P201 are available, which were, however, not considered as sensor data as input for the AD methods, but were only used to assign the concentrations measured by online NMR spectroscopy to the correct time step in the time series, cf. Supporting Information for details. In addition to the time-series data described above, several other data modalities were collected with the unconventional sensors described in Section 2.1. Operating-phase audio data are provided as a series of .wav files, each 40 seconds long. Images from the cameras described in Section 2.1 are compressed into .mp4 video files. Concentration data measured with online NMR and offline gas chromatography are supplied as raw tabular and processed time-series data. The supplied information enables assigning the point in time in the experiment to the data in all cases. Furthermore, for each experiment, tabular data sheets on the ambient conditions, the detailed plant control regime, and the log files from the PCS described in Section 2.4 are available. Moreover, general substance data for the respective systems, including residue curves and VLE behavior, are provided, along with tabular information on the initial mixture for each experiment.

4.2 Database

4.2.1 Structure

In the database, data of different modalities, i.e., different types of information, are stored. The database is structured into several levels, as shown in Figure 7. On the top level, general metadata on the plant is provided, including the information given in Section 2.1 of this paper. The following levels correspond to subdirectories in the database’s file structure. The actual data are found on the lowest level.

On the second level, the data are split according to their modality: times-series data of sensors, actuators, labels (+ anomaly metadata), audio, video, NMR, GC, and sensor uncertainties as well as metadata on the plant units, operation logs, ambient conditions, and

the substances and mixtures. The reason for splitting the data from the same experiment by their modalities at this high level is that, in many applications, data from certain modalities (e.g., audio, video) are not of interest. Furthermore, not all modalities are available for every experiment; an overview is given in the Supporting Information.

On the third level, different combinations (plant setups + systems) are distinguished. Setups are defined by function and describe the plant’s actual configuration, which varies only slightly – mainly when sensors or minor components are changed. At the fourth level, the operating point is specified, i.e., the set of all settings used in the plant operation, including parameters, controller setpoints, and fixed device power outputs. At the lowest level, the actual datasets are found, including both normal and anomalous runs. Anomalous runs are accompanied by metadata describing the anomaly and its cause, and both normal and anomalous experiments may include additional comments as metadata.

Hence, each individual dataset belongs to a directory tree along which all pertinent information on the experiment can be obtained. The tree structure also enables an unambiguous assignment of datasets of different modalities (e.g., sensor and actuator) to an experiment. This assignment is also facilitated by the filenames, which always contain information about the experiment.




Figure 7: Overview of the structure of the database created in the present work. The hierarchy level decreases from left to right. The colored folders indicate the classes on each level. The actual data are located on lowest level.

4.2.2 Datasets

Overview. In total, the database includes process data from 119 experiments, comprising more than one million time steps. Table 6 gives an overview of the number of experiments and operating points for each studied system. Table 7 summarizes the ranges of important parameters of the distillation experiments. The number of experiments often exceeds the number of operating points because, in many cases, especially when anomalies were introduced, the same operating point was chosen in different runs. Experiments with the system (ethanol + 2-propanol) were conducted during preliminary plant tests but are included in the database. Mixtures, substance purities, and residue-curve plots are provided in the Supporting Information and in the database as metadata. The time-series data are partitioned into start-up, operation, and shut-down phases; anomaly labels and metadata are available only for the operation phase.

Table 6: Overview of normal and anomalous experiments per system, including the number of experiments and unique operating points; operating points may repeat across experiments.

System	Type	No. of experiments	No. of operating points
1-Butanol + 2-propanol + water	Normal	29	25
	Anomalous	62	27
Acetone + 1-butanol + methanol	Normal	8	8
	Anomalous	16	12
Ethanol + 2-propanol	Normal	3	2
	Anomalous	1	1


Table 7: Overview of parameter ranges defining the operating point.

Parameter	Lower limit	Upper limit
Heat supply	105 W	510 W
Pressure	200 mbar	950 mbar
Reflux ratio	0.33	10

Examples. In Figure 8, data from an anomalous experiment is compared to data from the corresponding normal experiment. The figure shows results for the system (1-butanol + 2-propanol + water) and presents data for the temperatures along the column and the

pressure, as well as the composition of the mixture in the reboiler vessel V001 measured by NMR spectroscopy. Further sensor data for these experiments are provided in the Supporting Information. We start the discussion with the results from the normal experiment (upper two panels in Figure 8). During start-up, the temperature in the reboiler vessel (T703) increases rapidly and shows a kink when the boiling temperature of the feed is reached, as expected. The temperature measured below the first column section (T709) is similar to the temperature in the reboiler vessel (T703) throughout the experiment. The onset of evaporation also leads to a jump in the temperatures measured below the other column sections (T711, T712) and the temperature at the top of the column (T705). During the experiment, these temperatures remain relatively constant first, which is a result of the weak influence of the changes in the composition on the temperature. They begin to rise significantly when the mixture in the reboiler vessel V001 is almost depleted of water. At the end of the run, a small amount of highly purified 1-butanol, the heavy boiler in the studied system, remains in the reboiler vessel V001. The pressure is constant throughout the experiment.

The initial settings of the anomalous experiment (lower two panels in Figure 8) were the same as those for the normal experiment, which is why the signals obtained from both experiments are generally similar. There are two perturbations, which lead to anomalies and which can be discerned easily. The three phases of the perturbation (blind phase, anomalous phase, recovery phase) are marked in each case as described in Section 3.2. The perturbations consisted of changes in the setpoint of the heating H002 in the reboiler vessel V001 (normal setpoint: 105 W; first perturbation: temporary increase to 205 W; second perturbation: temporary decrease to 35 W).



4.3 Metadata of Anomalies

Contextual metadata is essential for understanding anomalies and ensuring transparency. While raw process signals capture responses to perturbations, metadata explain why these responses occur and how they relate to the underlying plant behavior. In addition to the experimental data, the database therefore includes structured metadata describing each labeled anomaly. To provide a standardized, semantically grounded representation, these metadata are formalized using an ontology. Semantic Web Technologies (SWTs)³⁹, with ontologies as a core component, offer a framework for capturing and organizing such knowledge in a uniform, machine-interpretable way. Ontologies define classes and relationships that represent domain knowledge; by instantiating these with concrete details about anomalies and their causes, the metadata become both consistent and interoperable.

The description of perturbations and their effects in this database follows the principles of Failure Mode and Effects Analysis (FMEA)⁴⁰. Here, the term failure is used interchangeably with perturbation, referring to a disturbance of the process. The adapted FMEA approach distinguishes between the perturbation mode (the manner in which a system is disturbed), the perturbation itself (the initiating event leading to that mode), and the perturbation effect (its immediate consequence)^{41,42}. The ontology applied in this work builds primarily on the Folio Ontology developed by Steenwinckel et al.⁴³, extended with concepts introduced by Klein et al.⁴⁴. The Folio Ontology aligns the FMEA perspective with the Semantic Sensor Network (SSN)⁴⁵ and the Sensor, Observation, Sample, and Actuator (SOSA) ontologies³⁶, allowing each perturbation mode to be linked both to the affected system component and to the sensor that records its observable effect. Klein et al.⁴⁴ further emphasize that a component’s perturbation mode is determined by its function.

The metadata are stored in YAML¹ format, providing a machine-interpretable yet human-readable representation. Figure 9 shows an example of such a YAML instance, describing the anomaly caused by a temporary pressure increase, cf. Figure 5. The example is explained

¹<https://yaml.org/>

in detail in the following.

```
1 anomalies:
2   - class: ConfirmedAnomaly
3     id: OP003TA002A001
4     hasBeginning: '14:55:31'
5     hasEnd: '14:58:00'
6     hasCategory: transient
7     hasRecoveryAction: Restore normal state
8     Perturbation: Increase in pressure setpoint PY23 by 200 mbar to 700 mbar
9     PerturbationMode:
10      label: Increased vacuum line throttling
11      hasBeginning: '14:55:05'
12      hasEnd: '14:58:00'
13      happenedAt:
14        System:
15          label: Vacuum Throttle Valve
16          id: TV1
17          hasFunction: Throttling of vacuum line for system pressure regulation
18          hasNormalState:
19            - automatic closed-loop controlled
20            - pressure setpoint 500 mbar
21            - sensor PY23
22          hasAnomalousState:
23            - automatic closed-loop controlled
24            - pressure setpoint 700 mbar
25            - sensor PY23
26      hasLocalEffect:
27        LocalEffect:
28          label: Increased system pressure
29          id: Observation001_OP003TA002A001
30          madeBySensor:
31            Sensor:
32              id: PY23
33          observedProperty:
34            ObservableProperty: Pressure
35          hasFeatureOfInterest:
36            FeatureOfInterest: Entire plant
```

Figure 9: Metadata stored in YAML format summarizing information on an anomaly caused by a temporary pressure increase.

In the example shown in Figure 9, the perturbation is a temporary increase of the system pressure setpoint. This perturbation disturbs the system's normal function and manifests as a specific perturbation mode, here, an increased flow resistance in the vacuum line caused

by the throttle valve within the pressure control loop. The throttle valve thus represents the affected system component. Operating in this mode produces a visible anomaly, which is linked to the corresponding sensor observations. Observations follow the SOSA pattern: each observation is made by a sensor, measures an observable property, and pertains to a feature of interest, i.e., the part of the plant whose property is being measured. In this example, the increased throttling of the vacuum line leads to a rise in system pressure, observed by a pressure sensor. Because this change affects the global pressure balance, the feature of interest is not limited to a specific component, such as a pipe segment or vessel, but encompasses the entire plant. A perturbation mode can also give rise to a local effect, representing a system-specific observable response that accompanies its onset. The concept of a local effect distinguishes these targeted, component-level observations from more general observations that may occur over time, also depending on the duration of the perturbation mode. Note that, in some cases, a perturbation may remain unobserved in the available sensor data. On the other hand, anomalies, describing observable deviations from normal plant operation, can be caused by both unexpected sensor deviations without an identifiable cause and be linked to a known perturbation that can be explicitly described (so-called confirmed anomalies); the above described example is a confirmed anomaly.

The ontology schema underlying the YAML instantiation is shown in Figure 10. Blue boxes denote ontology classes, labeled blue arrows indicate semantic relationships, and black arrows denote hierarchical “is a” relations. The directionality of blue arrows in the figure is limited to one side for clarity; however, the actual relations can be bidirectional. For example, the relation “hasCause” between a confirmed anomaly and a perturbation is conceptually complemented by the inverse relation “isCauseOf”.




Figure 10: Ontology schema with classes (boxes) and labeled relations (arrows). Blue arrows mark semantic relationships between classes and black arrows denote hierarchical “is a” relations.

In addition to semantic relations between the classes, those classes can also have descriptive attributes. A confirmed anomaly includes a recovery action, which describes how the system is restored to its normal condition. It also includes a category that denotes the temporal persistence of an anomaly and is classified as transient, recurrent, or permanent. Transient, thereby, refers to a continuous anomaly that resolves once the perturbation is removed. Recurrent indicates that a single perturbation causes an intermittent anomaly that alternates between observable and non-observable states over time. Permanent characterizes cases where the anomaly persists until the end of the experiment or does not recover even after the perturbation has been resolved. Both anomaly and perturbation mode include beginning and end attributes to record their temporal boundaries. For a perturbation mode during a confirmed anomaly, these timestamps cover the blind phase (called phase “1” in the label modality here) and anomalous phase (called phase “2” in the label modality), whereas for an observable anomaly, they span the anomalous phase and the recovery phase (called phase “3” in the label modality).

The system class includes attributes for its normal and anomalous states, characterizing fault-free and faulty conditions. Every instance of an anomaly, perturbation mode, sensor, observation, and system is identified by a globally unique ID. Sensor and system identifiers correspond to their names in the P&I diagrams; sensor identifiers also match the names of their data streams. An anomaly’s identifier is unique to that specific event. The attributes are summarized in Table 8.

Table 8: Class attributes introduced in the database.

Class	Attribute
Anomaly, Perturbation mode	<i>hasBeginning, hasEnd, id</i>
Confirmed anomaly	<i>hasCategory</i>
System	<i>hasNormalState, hasAnomalousState, id</i>
Observation, Sensor	<i>id</i>
Confirmed anomaly	<i>hasRecoveryAction</i>

Each ontology class is instantiated with plain-text values rather than a fixed hierarchy of predefined phrases. This deliberately lightweight design keeps the representation flexible while remaining semantically grounded. For the database presented here, domain experts annotated each anomaly according to the ontology’s structure and through careful visual inspection of the corresponding sensor readings.

Each YAML file contains the anomaly instances of a single experiment. Anomalies are listed in chronological order, and filenames encode both the operating point and the anomalous experiment. As an example, `OP003TA002A001` identifies the first anomaly (001) associated with the second time series for operating point 3 (file: `operating_point_003/test_anormal_experiment_002.csv`). Observation IDs are sequential and reference their associated anomaly; for instance, `Observation001_OP003TA002A001` denotes the first observation linked to anomaly `OP003TA002A001`. When multiple components were perturbed simultaneously, each component exhibited its own perturbation mode. However, these modes are grouped under a single anomaly ID representing the shared perturbation, e.g., a simultaneous shutdown of all mantle heaters. Currently, only the local effect of each mode is documented, describing its immediate, component-specific impact. This structure can

be readily expanded to capture more complex relationships, such as causal chains and the propagation of anomalies through the system.

This explicit mapping renders the metadata both machine-interpretable and human-readable. Each anomaly thus becomes not only a labeled segment of sensor data but also a coherent knowledge unit describing what happened, why it happened, and how it was observed. It therefore provides a valuable foundation for future applications, particularly for developing AD methods that not only identify anomalies but also infer their underlying causes.

5 Conclusions

This work closes a major gap in the open literature that has long hindered the development of machine learning (ML) methods for anomaly detection (AD) in chemical processes: the absence of suitable experimental data. Until now, only a single synthetic dataset from process simulations has been publicly available. Industrial data remain inaccessible due to confidentiality and, even if available, would lack consistent annotation. We therefore provide the first extensive experimental database specifically designed to support the development of ML-based AD methods in chemical process systems.

The data originate from a laboratory-scale batch distillation plant, specifically designed and equipped for the present study, and operated under a wide range of conditions and mixtures, including both fault-free and anomalous experiments. Alongside conventional process variables, the database includes multimodal information, such as concentration data from online NMR spectroscopy, as well as video and audio recordings. Expert annotations describe the anomalies and their causes based on an ontology developed in this work. The complete database is publicly available in a structured format under a CC BY 4.0 license at doi.org/10.5281/zenodo.17395544.

The studied batch distillation process is inherently unsteady. Complementary data from a steady-state continuous distillation process are currently being collected at TU Munich,

within the same coordinated project, the DFG Research Unit 5359 “Deep Learning on Sparse Chemical Process Data.” These data will be released in a format consistent with the present database.

Future work will extend the database to additional operating modes of the batch distillation plant, including dynamic control and fed-batch operation. This will enable transferability studies and the evaluation of cross-system generalization in ML-based AD methods. Expanding the chemical space to multicomponent and liquid–liquid systems – potentially involving plant modifications such as a decanter – will further broaden the database’s scope. Beyond serving as a benchmark for ML-based AD, the database provides a foundation for future research. Promising directions include interpretable and explainable ML models that diagnose anomalies, as well as data-driven strategies for anomaly mitigation and control. The multimodal nature of the data allows studies on sensor fusion, temporal alignment, and uncertainty-aware learning. Moreover, the ontology and metadata foster standardized representations of process anomalies, enhancing transferability and comparability across systems. The database also supports training of generative AI models that transfer simulation data into realistic synthetic data, complementing experimental databases for ML-based AD training. Overall, the database bridges the gap between realistic experimental data and the methodological needs of advanced ML research. It is expected to promote collaboration between the chemical engineering and data science communities and to accelerate progress toward more reliable, transparent, and autonomous process operation.

Acknowledgements

The authors gratefully acknowledge funding from the Deutsche Forschungsgemeinschaft (DFG), as this work was conducted within the DFG Research Unit FOR 5359 "Deep Learning on Sparse Chemical Process Data" (grant number 459419731). Moreover, we thank the DFG Gerätezentrum "Laboratory for Advanced Spin Engineering – Magnetic Resonance (LASE-

MR)" (grant number 537627671) for its technical support.

Supporting Information

Supporting information is available for this work and includes:

- Additional information on the batch distillation plant
- Information on sensor calibration
- Information on ambient data collection
- Information on offline and online concentration measurement
- Additional information on chemical compounds and mixtures
- Additional example sensor data

References

- (1) Chandola, V.; Banerjee, A.; Kumar, V. Anomaly detection: A survey. *ACM Computing Surveys* **2009**, *41*, 1–58.
- (2) Venkatasubramanian, V.; Rengaswamy, R.; Yin, K.; Kavuri, S. N. A review of process fault detection and diagnosis Part I: Quantitative model-based methods. *Computers & Chemical Engineering* **2003**, *27*, 293–311.
- (3) Venkatasubramanian, V.; Rengaswamy, R.; Kavuri, S. N. A review of process fault detection and diagnosis Part II: Qualitative models and search strategies. *Computers & Chemical Engineering* **2003**, *27*, 313–326.
- (4) Venkatasubramanian, V.; Rengaswamy, R.; Kavuri, S. N.; Yin, K. A review of process fault detection and diagnosis Part III: Process history based methods. *Computers & Chemical Engineering* **2003**, *27*, 327–346.
- (5) Chiang, L. H. In *Fault Detection and Diagnosis in Industrial Systems*; Russell, E. L., Braatz, R. D., Eds.; Advanced Textbooks in Control and Signal Processing; Springer: London, 2001.
- (6) Hodge, V.; Austin, J. A Survey of Outlier Detection Methodologies. *Artificial Intelligence Review* **2004**, *22*, 85–126.
- (7) Mowbray, M.; Vallerio, M.; Perez-Galvan, C.; Zhang, D.; Del Rio Chanona, A.; Navarro-Brull, F. J. Industrial data science – a review of machine learning applications for chemical and process industries. *Reaction Chemistry & Engineering* **2022**, *7*, 1471–1509.
- (8) Dobbelaere, M. R.; Plehiers, P. P.; Van de Vijver, R.; Stevens, C. V.; Van Geem, K. M. Machine Learning in Chemical Engineering: Strengths, Weaknesses, Opportunities, and Threats. *Engineering* **2021**, *7*, 1201–1211.

- (9) Göttl, Q.; Pirnay, J.; Burger, J.; Grimm, D. G. Deep reinforcement learning enables conceptual design of processes for separating azeotropic mixtures without prior knowledge. *Computers & Chemical Engineering* **2025**, *194*, 108975.
- (10) Gond, D.; Sohns, J.-T.; Leitte, H.; Hasse, H.; Jirasek, F. Hierarchical matrix completion for the prediction of properties of binary mixtures. *Computers & Chemical Engineering* **2025**, *199*, 109122.
- (11) Hayer, N.; Jirasek, F.; Hasse, H. Prediction of Henry’s law constants by matrix completion. *AIChE Journal* **2022**, *68*.
- (12) Hayer, N.; Hasse, H.; Jirasek, F. Modified UNIFAC 2.0-A Group-Contribution Method Completed with Machine Learning. *Industrial & Engineering Chemistry Research* **2025**, *64*, 10304–10313.
- (13) Hayer, N.; Specht, T.; Arweiler, J.; Hasse, H.; Jirasek, F. Similarity-Informed Matrix Completion Method for Predicting Activity Coefficients. *The Journal of Physical Chemistry A* **2025**, *129*, 3141–3147.
- (14) Hayer, N.; Wendel, T.; Mandt, S.; Hasse, H.; Jirasek, F. Advancing thermodynamic group-contribution methods by machine learning: UNIFAC 2.0. *Chemical Engineering Journal* **2025**, *504*, 158667.
- (15) Hoffmann, M.; Hasse, H.; Jirasek, F. GRAPPA – A Hybrid Graph Neural Network for Predicting Pure Component Vapor Pressures. 2025.
- (16) Jirasek, F.; Hasse, H. Combining Machine Learning with Physical Knowledge in Thermodynamic Modeling of Fluid Mixtures. *Annual Review of Chemical and Biomolecular Engineering* **2023**, *14*, 31–51.
- (17) Jirasek, F.; Alves, R. A. S.; Damay, J.; Vandermeulen, R. A.; Bamler, R.; Bortz, M.; Mandt, S.; Kloft, M.; Hasse, H. Machine Learning in Thermodynamics: Prediction of

- Activity Coefficients by Matrix Completion. *The Journal of Physical Chemistry Letters* **2020**, *11*, 981–985.
- (18) Jirasek, F.; Hasse, H. Perspective: Machine Learning of Thermophysical Properties. *Fluid Phase Equilibria* **2021**, *549*, 113206.
- (19) Specht, T.; Arweiler, J.; Stüber, J.; Münnemann, K.; Hasse, H.; Jirasek, F. Automated nuclear magnetic resonance fingerprinting of mixtures. *Magnetic Resonance in Chemistry* **2023**, *62*, 286–297.
- (20) Specht, T.; Nagda, M.; Fellenz, S.; Mandt, S.; Hasse, H.; Jirasek, F. HANNA: hard-constraint neural network for consistent activity coefficient prediction. *Chemical Science* **2024**, *15*, 19777–19786.
- (21) Vollmer, L.; Fellenz, S.; Jirasek, F.; Leitte, H.; Hasse, H. KnowTD—An Actionable Knowledge Representation System for Thermodynamics. *Journal of Chemical Information and Modeling* **2024**, *64*, 5878–5887.
- (22) Chadha, G. S.; Rabbani, A.; Schwung, A. Comparison of Semi-supervised Deep Neural Networks for Anomaly Detection in Industrial Processes. 2019 IEEE 17th International Conference on Industrial Informatics (INDIN). 2019; pp 214–219.
- (23) Monroy, I.; Escudero, G.; Graells, M. *19th European Symposium on Computer Aided Process Engineering*; Elsevier, 2009; pp 255–260.
- (24) Inoue, J.; Yamagata, Y.; Chen, Y.; Poskitt, C. M.; Sun, J. Anomaly Detection for a Water Treatment System Using Unsupervised Machine Learning. 2017 IEEE International Conference on Data Mining Workshops (ICDMW). 2017; pp 1058–1065.
- (25) Song, B.; Suh, Y. Narrative texts-based anomaly detection using accident report documents: The case of chemical process safety. *Journal of Loss Prevention in the Process Industries* **2019**, *57*, 47–54.

- (26) Tian, W.; Liu, Z.; Li, L.; Zhang, S.; Li, C. Identification of abnormal conditions in high-dimensional chemical process based on feature selection and deep learning. *Chinese Journal of Chemical Engineering* **2020**, *28*, 1875–1883.
- (27) Wu, G.; Zhang, Y.; Deng, L.; Zhang, J.; Chai, T. Cross-Modal Learning for Anomaly Detection in Complex Industrial Process: Methodology and Benchmark. 2024.
- (28) Hartung, F. et al. Deep Anomaly Detection on Tennessee Eastman Process Data. *Chemie Ingenieur Technik* **2023**, *95*, 1077–1082.
- (29) Russell, E. L.; Chiang, L. H.; Braatz, R. D. *Data-driven Methods for Fault Detection and Diagnosis in Chemical Processes*; Springer London, 2000.
- (30) Schmidl, S.; Wenig, P.; Papenbrock, T. Anomaly detection in time series: a comprehensive evaluation. *Proceedings of the VLDB Endowment* **2022**, *15*, 1779–1797.
- (31) Zamanzadeh Darban, Z.; Webb, G. I.; Pan, S.; Aggarwal, C.; Salehi, M. Deep Learning for Time Series Anomaly Detection: A Survey. *ACM Computing Surveys* **2024**, *57*, 1–42.
- (32) Wang, F.; Jiang, Y.; Zhang, R.; Wei, A.; Xie, J.; Pang, X. A Survey of Deep Anomaly Detection in Multivariate Time Series: Taxonomy, Applications, and Directions. *Sensors* **2025**, *25*, 190.
- (33) Downs, J.; Vogel, E. A plant-wide industrial process control problem. *Computers & Chemical Engineering* **1993**, *17*, 245–255.
- (34) Rieth, C. A.; Amsel, B. D.; Tran, R.; Cook, M. B. Additional Tennessee Eastman Process Simulation Data for Anomaly Detection Evaluation. 2017.
- (35) Kister, H. What Caused Tower Malfunctions in the Last 50 Years? *Chemical Engineering Research and Design* **2003**, *81*, 5–26.

- (36) Janowicz, K.; Haller, A.; Cox, S. J.; Le Phuoc, D.; Lefrançois, M. SOSA: A lightweight ontology for sensors, observations, samples, and actuators. *Journal of Web Semantics* **2019**, *56*, 1–10.
- (37) Gronle, M.; Lyda, W.; Wilke, M.; Kohler, C.; Osten, W. itom: an open source metrology, automation, and data evaluation software. *Applied Optics* **2014**, *53*, 2974.
- (38) CARPi, N.; Minges, A.; Piel, M. eLabFTW: An open source laboratory notebook for research labs. *The Journal of Open Source Software* **2017**, *2*, 146.
- (39) Ye, J.; Dasiopoulou, S.; Stevenson, G.; Meditskos, G.; Kontopoulos, E.; Kompatsiaris, I.; Dobson, S. Semantic web technologies in pervasive computing: A survey and research roadmap. *Pervasive and Mobile Computing* **2015**, *23*, 1–25.
- (40) Burge, S. The Systems Engineering Tool Box. 2010; <https://www.burgehugheswalsh.co.uk/Uploaded/1/Documents/FFMEA-Tool-v2.pdf>.
- (41) Müller, C.; Lunde, R.; Hönig, P. Generation of a Failure Mode and Effects Analysis with SmartIflow. Proceedings of the 30th European Safety and Reliability Conference and 15th Probabilistic Safety Assessment and Management Conference. 2020; pp 1662–1669.
- (42) Pecht, M.; Gu, J. Physics-of-failure-based prognostics for electronic products. *Transactions of the Institute of Measurement and Control* **2009**, *31*, 309–322.
- (43) Steenwinckel, B.; Heyvaert, P.; De Paepe, D.; Janssens, O.; Vanden Haute, S.; Dimou, A.; De Turck, F.; Van Hoecke, S.; Ongenaes, F. Towards adaptive anomaly detection and root cause analysis by automated extraction of knowledge from risk analyses. Proceedings of the 9th international semantic sensor networks workshop, co-located with 17th international semantic web conference (ISWC 2018). 2018; pp 17–31.

- (44) Klein, P.; Malburg, L.; Bergmann, R. Combining informed data-driven anomaly detection with knowledge graphs for root cause analysis in predictive maintenance. *Engineering Applications of Artificial Intelligence* **2025**, *145*, 110152.
- (45) Compton, M. et al. The SSN ontology of the W3C semantic sensor network incubator group. *Journal of Web Semantics* **2012**, *17*, 25–32.