An ontology for maintenance activities and its application to data quality

Caitlin Woods [a,*], Matt Selway [b], Tyler Bikaun [a], Markus Stumptner [b] and Melinda Hodkiewicz [c]

[a] *Department of Computer Science and Software Engineering, University of Western Australia, Crawley, 6009, Australia*
E-mails: caitlin.woods@research.uwa.edu.au, tyler.bikaun@research.uwa.edu.au
[b] *Industrial AI Research Centre, UniSA STEM, University of South Australia, Adelaide, 5000, SA, Australia*
E-mails: matt.selway@unisa.edu.au, markus.stumptner@unisa.edu.au
[c] *School of Engineering University of Western Australia, Crawley, 6009, Australia*
E-mail: melinda.hodkiewicz@uwa.edu.au

Abstract.

Maintenance of assets is a multi-million dollar cost each year for asset intensive organisations in the defence, manufacturing, resource and infrastructure sectors. These costs are tracked though maintenance work order (MWO) records as a MWO is generated every time work is done on an asset. MWO records are the only digital record of maintenance events for the life history of an asset and hundreds of thousands are generated every year in asset intensive organisations. The information contained in MWOs is necessary for determining what maintenance work is done when, to what item, and how much it cost. This knowledge is used in many downstream data processing tasks such as assessing compliance with maintenance strategy, root cause analysis, budget compliance, and business process improvement. Hence, data quality is important.

MWO records contain structured data for dates, costs, and asset identification but a key field of interest contains unstructured text created by operators and maintenance technicians to describe the work required, for example ‘replace leaking pump’. Our focus in this paper is on data quality of the maintenance activity term e.g. replace, repair, adjust and inspect. The data quality use case is to assess if a maintenance activity described in unstructured MWO text is consistent with the structured data contained in the MWO record.

We present two contributions in this paper. First, we propose a reference ontology for maintenance activity terms. We use natural language processing and a manual clustering exercise to identify seven core maintenance activity terms and their synonyms from 21,088 MWOs. We provide elucidations for these seven terms. Second, we demonstrate use of the reference ontology in an application-level ontology using an industrial use case. The end-to-end NLP-ontology pipeline identifies data quality issues with 55% of the MWO records for a centrifugal pump over 8 years. For 22% of records the ontology can infer the activity when it was not provided in the unstructured text and for 33% of records it identifies inconsistencies between the activity term used and the other information in the record and routes these to an engineer for checking.

Alignment to International Standards is an important consideration for many industry sectors and the selection of the maintenance activity terms is informed by the ISO 14224 and ISO 15926-4 standards and conforms to ISO/IEC 21838-2 Basic Formal Ontology (BFO). The reference and application ontologies presented here provide an example for how industrial organisations can augment their maintenance work management processes with ontological workflows, specifically to improve data quality.

Keywords: maintenance work order, ontology, natural language processing

[*]Corresponding author. E-mail: caitlin.woods@research.uwa.edu.au.

1. Introduction

Maintenance of assets is a significant cost input for the manufacturing, resources, defence and infrastructure sectors. Maintenance costs typically range between 20–60% of operational expenditure depending on industry and asset type [1]. There has been significant effort in the last decade to move from reactive to preventative and predictive maintenance strategies propelled by developments in sensing, WiFi, cloud computing, and data analytics. However generating value from analytics using these platforms has often proved challenging [2–4]. In part, this is due to the way in which data describing what maintenance work was actually done, to what item, when it was done and what it cost, is captured and stored [5, 6] in maintenance work order (MWO) records; that is, as unstructured free text.

MWOs are the equipment equivalent of health records for humans. Tens to hundreds of thousands of MWOs are generated a year on a moderately complex site or asset system. A MWO is an artifact that is generated in industrial organisations (that are mature enough to maintain a maintenance strategy) to inform technicians that work needs to be done. MWOs can be generated automatically (by a Computerised Maintenance Management System) in the case of routine maintenance, or manually by a technician. These MWOs contain an unstructured text field with a short description of the job (e.g., "replace all damaged idlers" or "inspect / refurb scrapers"). Each record also has structured data fields for a unique work order number, asset functional location, dates for the proposed and actual start and end date of the work, budget and actual costs, as well as fields for metadata about the record such as when it was generated and closed. Computer readability of the unstructured text in MWOs using natural language processing (NLP) is currently an active area of research and development [7, 8].

Our focus in this paper is on data quality of the maintenance activity term in the MWO. This activity is described by verbs such as replace, repair, adjust and inspect. This verb is written manually into the unstructured short text of the record by an operator or maintainer. If this information is unclear or unreliable it adds uncertainty to the identification of life cycle events such as end-of-life and when preventative actions were done (or not). While it desirable to predict a failure event from condition monitoring data, the engineer also needs to know if action was taken based on the prediction, what the maintenance activity was, and if the action was taken before the failure manifest. The raw data from which this action can be inferred is stored as structured and unstructured text in maintenance work order records (MWOs). Currently this inference is done manually by engineers and planners.

Data quality of MWOs has been a topic of organisational behaviour research for many years [9–11]. However the problem is unresolved and data quality is cited as an issue affecting maintenance and warranty management improvement programs [12–14]. Labour productivity in maintenance is now a focus of senior management attention [1]. The use of experienced maintenance planner and reliability engineer time to manually read and process MWOs to get information for their analysis is a significant impediment to improving their productivity [15]. Machine augmentation of this basic task is required.

Our interest is improving data quality using ontology-based reasoning on MWO unstructured texts pre-processed with natural language processing techniques. Currently questions such as ‘was pump 101 actually replaced on 12/1/20?’ are resolved manually by a reliability engineer [16, 17]. Such analysis requires the reliability engineer to examine MWO records one at a time. However, the reasoning used to answer a number of these questions can be formalised into a set of rules [18]. This opens the door to the use of ontologies as part of the MWO data quality improvement process. The specific focus of this paper is an ontology for maintenance activity terms.

We propose a solution to improve MWO data quality using a reference ontology and an application-level ontology. The reference ontology contains a holistic view of the types of activities conducted by maintenance personnel. Elucidations based on ontological analysis are provided for seven core maintenance activity terms. The application-level ontology uses the reference ontology and real-world MWO data to perform the necessary reasoning tasks for our use case. This ontology examines each MWO and determines if the maintenance activity word used in the MWO is consistent with the other information in the maintenance record.

If the activity described in a MWO’s unstructured text matches the record’s structured data, then we can be more certain that the activity as described actually occurred in practice. This information can then be used, with confidence, for a number of important downstream analytics tasks such as in calculating reliability metrics such as mean-time-to-failure, in root cause analysis following major undesirable events, and for assessing the effectiveness of maintenance strategy [12, 19–21].

The paper is organised as follows. Section 2 describes previous ontology development and industry reference data models in the maintenance area and provides more details on MWO and their contents. The use case including instance data are presented in Section 3. Section 4 describes steps in the process of identifying the 7 maintenance activity classes and their elucidations. Ontology design and competency questions are presented in Section 5 and the performance of the ontology is assessed in the evaluation in Section 6. Finally, Section 7 critically examines the performance of the ontology and identifies opportunities for further work.

2. Background

2.1. Maintenance

Maintenance is defined by an International Standard as "the actions intended to retain an item in, or restore it to, a state in which it can perform a required function" [22]. We note from this that the notion of an action or activity is central to the concept of maintenance.

In mature maintenance organisations each maintainable item has an associated maintenance strategy. Maintenance strategy determines what work should be done, when to do it and at what level in the asset hierarchy it should be performed. Deciding on maintenance strategy is a well established process based on failure modes and effects analysis and reliability centred maintenance (RCM). These are described in international standards such as [23, 24]. The 'what strategy?' decision is informed by factors such as the safety, environmental, production or cost consequences of an item's failure, if loss of function is technically and cost effective to observe and how long it takes from observation of deterioration to failure event. RCM decision-logic guides the user towards one of the following strategies: a) use based, also known as fixed interval restoration/ repair/ inspect strategy, b) strategy based on condition monitoring and inspections, c) failure-finding and d) run-to-failure maintenance strategies [25]. Colloquially these are classified into preventative and corrective maintenance strategies, with *preventative* including fixed interval work and periodic inspections such as condition monitoring and *corrective* covering failure finding, run to failure and work arising from preventative tasks. The notions of preventative and corrective strategies are relevant to the reasoning used in this paper.

2.2. Review of previous ontology work on maintenance activity terms

There is increasing interest in the being able to automatically process equipment maintenance and failure data using ontologies as a means of standardising decisions made in processing this data for decision support, and automating tedious data processing tasks that require the input of experienced people [12, 15, 20, 21]. Interest is high in industries such as defence, aerospace, oil and gas, and manufacturing and this is compelling groups to come together to consider standardisation as they recognise the scale of the task and the challenge of going it alone. In 2021 the ISO/IEC 21838-1 for Top Level Ontologies [26] was issued alongside ISO/IEC 21838-2 Standard [27] describing Basic Formal Ontology (BFO). Standardisation of a further two top level ontologies, Descriptive ontology for linguistic and cognitive engineering (DOLCE) and TUpper is in progress. This standardisation process is a key plank in industry acceptance of ontologies. There is a surge of activity to develop industry-relevant, domain specific ontologies, that are aligned to one of these Top Level Ontologies. Of relevance to this paper is the work of the Industrial Ontology Foundry (IOF) and their work to publish a domain ontology for manufacturing aligned to BFO. The IOF promotes a principles-based approach to the design of ontologies for design, maintenance, supply chain, production, and lifecycle management of equipment. A number of the terms and relations necessary to support description of maintenance activities are already formalised in IOF [28].

While there have been a number of papers published on maintenance ontologies these have, in the main, been developed using a top down approach. Some have tried to capture maintenance in general [29, 30] while others have focused more specifically on various processes in maintenance such as maintenance work management [31], failure modes and effects analysis [32], fault classification from warranty records [33]. Dealing with real life maintenance data has not been the primary motivation for ontology development. Instead conceptual models have dominated with limited use cases selected to support and illustrate specific reasoning and conceptualisation challenges. There have

Table 1
Examples of other fields of interest in MWO records

Field(s)	**Description**	Field type
Functional location	Asset location where maintenance is performed	Structured text
Work description	Description of the problem or work to be done	Unstructured text
Work order number	Unique identifier for the record	Numeric
Work order type	Code to describe if the work is corrective or preventative	Code
Start and end dates (MWO)	Metadata for the MWO record	Date
Scheduled start and end dates	Scheduled start and end date for the proposed maintenance activity	Date
Actual start and end dates	Actual start and end date for the proposed maintenance activity	Date
Estimated and actual labour hours	Budget and actual manpower hours for the maintenance activity	Numeric
Estimated and actual labour costs	Budget and actual manpower cost for the maintenance activity	Numeric
Actual parts costs	Cost of parts used in the maintenance activity	Numeric

been some exceptions. For example, in the automotive industry Rajpathak and General Motors, have used natural language processing and ontologies to extract data from hundred of thousands of unstructured warranty records [20, 33, 34] and their is similar interest in the analysis of safety data from accident databases [35]. In the oil and gas industry, there is emerging work on using ontology patterns to ingest instance data from MWO's and failure modes and effects analysis records into ontologies for use in reasoning about the effectiveness of maintenance strategy for a complex operating asset [36].

When looking for previous ontology work specifically on maintenance activity classes we located a list of 16 drawn from a data set of 654 activities/ tasks for a wire harness assembly case study [37]. However these activities are defined using natural language phrases such as 'laying cable flat' and 'inserting into the tube or sleeve' and can be used only in this specific use case. The activity classes identified are not generic to maintenance in general.

2.3. Maintenance work order records

MWO is generated every time work is done on an asset. MWO records are the only digital record of maintenance events for the life history of an asset. Hundreds of thousands of these MWOs are generated every year in asset intensive organisations. In a MWO record the primary interest to the reliability engineers are the fields described in Table 1. The functional location is the asset identifier, the work order description contains the unstructured text of interest in our work, the work order type identifies if the work is preventative or corrective, and there are fields for the start and end dates and various budget and actual costs. Other fields, not shown, include meta data for the record. The information contained in MWOs is the only record of what maintenance work is done when, to what item, and how much it cost.

2.4. Natural language processing of maintenance work orders

The reader will note from Table 1 that there is no dedicated fields for key information such as "what specific item needs work?", "what work needs to be done?", "why does it need to be done?". The answers to these questions have to be inferred from a 4-8 word sentence captured in the 'Work description' field. A typical example, drawn from [38] is 'change out leaking engine'. This phrase is illustrative of MWOs in general in that it contains an activity word change out, a state or problem word 'leaking' and an item identifier 'engine'. Work to develop pipelines to separate the words in these MWOs into classes using natural language processing is very active [7, 20, 21, 38–40].

Once extracted, the information in these activity, item and state fields is being used for visualisation and analysis to support a number of maintenance and reliability decisions [16, 18, 33, 41]. However one challenge for all users is the large number of terms in all these classes resulting from synonyms, misspellings, different tenses, abbreviations, and jargon.

Work on defining classes for types of maintenance *activity* in ontology is limited. There is a list of 10 maintenance action codes (inspection, servicing, adjustment, alignment, resetting, replenishing, coating, repair, replace, overhaul) produced by [42] but the selections on the list are not justified or the terms fully described. A Bag-Of-Words approach on a corpus of 690,000 MWOs identified 103 maintenance terms but acknowledged that many terms such as service that can be a noun or a verb were missed [38]. There was no attempt to group these terms into manageable clusters. The raw data arising from these works are not freely available and when there are lists of classes they are not clearly defined. So instead we turn to lists available in the engineering standards literature.

2.5. Lists of maintenance activity terms from international Standards

We draw on two international standards, ISO 14224 (2016) and ISO 15926-4, for their lists of maintenance activity terms as follows.

2.5.1. ISO 14224

The ISO 14224 standard (Petroleum, petrochemical and natural gas industries - Collection and exchange of reliability and maintenance data for equipment) [43] is an influential international standard originally developed for the oil and gas industry but now widely used in other process and heavy industry sectors. This standard was originally developed in the early 1980s due to work by the oil and gas sector seeking to standardise maintenance and failure data collection for an Offshore Reliability Data handbook called OREDA [45]. Data quality is a key consideration in OREDA because the data in this handbook is so widely used by industry for safety-critical design decisions [46]. ISO 14224 has been through 5 editions and contains widely-used lists for equipment hierarchies, failure modes, mechanisms and causes and a set of maintenance activity terms. We list the activity terms and their definitions from ISO 14224 in Table 2.

2.5.2. ISO 15926-4 Activity class hierarchy

The ISO/TS 15926-4 Standard (Industrial automation systems and integration — Integration of life-cycle data for process plants including oil and gas production facilities — Part 4: Initial reference data) is a data model developed for the engineering design community particularly in the process sector. It includes a number of reference data models for activities, rotating equipment, static equipment, instrumentation, units of measure, valves etc. In the activity data model there are 1795 activity terms organised into a deep taxonomy (7 levels) from a superclass set of 22 terms. These terms are: acting, absorbing, affirming, behaving, being, choosing, competing, consigning, directing, disjoining, drawing, happening, occurring, passing-letting, process, pushing, reacting, reducing, refusing, relieving, squeezing, watching. Maintenance activity terms relating specifically to interactions with equipment such as replacing, repairing, inspecting sit under the superclass *acting* with the exception of troubleshooting and maintaining which are classified under *being*. A summary of the maintenance-related terms and their subclass structure is shown in Table 3. Descriptions of these terms are in Table 2.

2.5.3. Review of existing activity terms

Table 2 shows that there is agreement between the maintenance activity lists from ISO 14224 and ISO 15926 for the terms replace, repair, modify, adjust, check, service, test, inspect and overhaul. The descriptions in ISO 14224 incorporate information of potential value to reasoning, namely if the activity is corrective, preventative or can be both, as this information is available in the 'work type' field of the MWO. Activities associated only with preventative work also have the word 'periodic' in the description.

3. Use case

The use case for this ontology involves real-world industrial data. Table 4 contains a set of MWOs for a single functional location for a single pump in a process plant from 2012-2020.

Table 2

Maintenance activity terms and associated descriptions from ISO14224 [43] and ISO 15926-4 [44]

Maintenance Activity Term	**ISO 14224 descriptions.** The (C, P) indicates corrective and preventative maintenance strategies.	**ISO 15926-4 descriptions**
Replace	Replacement of an item by a new or refurbished item of the same type and make (C,P).	Acting by putting something new in the place of.
Repair	Manual maintenance action performed to restore an item to its original appearance or state (C).	Restoring by bringing an item back in its original state (shape and properties) or in a revised design state.
Modify	Replace, renew or change the item, or a part of it, with an item/ part of a different type, make, material or design (C, P).	Acting by replacing, renewing, or changing items, or parts with another item or part of different type, make, material or design.
Adjust	Bringing any out of tolerance condition into tolerance (C,P).	Changing by bringing in a more satisfactory state.
Refit	Minor repair/ servicing activity to bring back an item to an acceptable appearance, internal or external (C, P).	Not defined in ISO 15926-4
Check	The cause of failure is investigated but no maintenance action is performed, or action is deferred. Able to restart by simple actions e.g.restart or resetting (C).	Verifying by investigating and obtaining confirmation or substantiation of accuracy, fitness, or due performance
Service	Periodic service tasks: normally not dismantling of the item (e.g. cleaning, oil top up etc) (P).	Not defined in ISO 15926-4 from a maintenance task perspective.
Test	Periodic test of function or performance (P).	Validating that a physical object satisfies specified criteria; typically by verification of a property value of the physical object.
Inspect	Periodic inspection/ check: a careful scrutiny of an item carried out with or without dismantling, normally by the use of senses. (P).	Acting by viewing closely and critically in order to ascertain quality or state, detect errors, or otherwise appraise
Overhaul	Major overhaul (C,P).	Repairing as to restore to satisfactory working order.
Calibrate	Not defined in ISO 14224	Adjusting precisely for a particular function, or standardising by determining the deviation from standard and adjusting if appropriate.
Maintain	Not defined in ISO 14224	Keeping in a state of repair, efficiency, or validity, preserving from failure.
Troubleshoot	Not defined in ISO 14224	Maintaining by locating trouble and making repairs in machinery and technical equipment.
Renew	Not defined in ISO 14224	Revising by replacing old parts with new ones.

There are hundreds of other similar pumps in this facility. Columns 2-3 and 6-8 in Table 4 are representative of a data set a reliability engineer would normally work with. These fields are usually of most interest to the reliability engineer and are often the most readily available as reports downloadable as CSV files from the computerised maintenance management system (CMMS). The first column contains an ID to assist us with the competency questions. The second column contains a date, in this case we have selected the date on which the work was started. The third column was described earlier and contains the unstructured text describing the work. The sixth column describes the work order type as either corrective or preventative and the final two columns contain the labour and material costs.

Table 4 contains two columns not found in the CMMS. The fourth and fifth are calculated columns that capture the outputs of processing the unstructured text with an NLP pipeline. The ability to do this depends on lexical normalisation [47] and annotation of a large corpus of MWOs [40] which has been used to train a deep learning model. The outputs from this processing that are of interest to the work in this paper are the the extraction of the *item* and *activity* in the unstructured text. For each NLP-identified *item*, we have also provided the item's subunit. This subunit has been extracted from Table 5. The NLP pipeline that enables the work presented in this paper is discussed in further detail in Section 4.1.1.

Table 3
Mapping of activity super-class terms from ISO/TS 15926-4 (Industrial automation systems and integration — Integration of life-cycle data for process plants including oil and gas production facilities — Part 4: Initial reference data) showing their maintenance-related subclass terms.

Activity super-class	Text description (ISO 15926-4)	Maintenance-related subclasses from ISO 15926-4)
acting	Activity of carrying out a process of change or alteration	replacing → *acting* overhauling → repairing → restoring → *putting* → *acting* renewing → *revising* → *improving* → *enhancing* → *increasing* → *making* → *acting* testing → *validating* → *providing* → *supplying* → *acting* service → *service* − *acting* → *acting* inspecting → *acting* cleaning → *removing* → *acting* calibrating → adjusting → *changing* → *making* → *acting* modifying → *acting* analyzing → *studying* → *acting* greasing → lubricating → *applying* → *using* → *acting*
being	Activity that entails the having of an objective existence	troubleshooting → maintaining → *keeping* → *holding-having* → *having* → *being*

Over the 7.5 year period there have been 36 MWOs raised against the pumping system in this functional location. The majority (16 out of 36) are related to the pressure switch. The pressure switch is part of the Control and Monitoring subunit. The pump has been replaced once, in 2017, at a material cost of $> \$15,000$. A costly activity was performed on the pump housing due to an oil leak in 2018 for $\approx \$8,000$.

The MWOs include both preventative and corrective work, with preventative work being associated with a semi-structured form in unstructured text field usually starting with some periodic interval such as 78W for 78 weeks. Preventative work passes through the planning and scheduling process in the maintenance work management system [48]. Corrective work orders are urgent and need to be actioned immediately by shift maintainers whereas correctives go into the maintenance management process to be dealt with by maintenance planners and then scheduled into a weekly maintenance plan. In general there is a desire to move towards maintainers spending their time on preventative work with correctives generated from preventative inspections. Correctives are costly and disruptive and typically represent a failure of strategy, except where they associate with a deliberate run-to-failure strategy.

Finally of note is that some MWOs (e.g. 6, 12, and 19 in Table 4) have $\$0$ labour cost, which is interpreted by reliability engineers as unexecuted work. This frequently occurs as work can be scheduled but not actioned due to production, resourcing, or other issues. Consequently, it is important to be track work that is not executed and we use this information in our reasoning.

There are many data models to describe how asset systems can be decomposed into functional or physical breakdown. A detailed review of these is beyond the scope of this paper. Our taxonomy is informed by the hierarchy in ISO 14224 [43] of maintainable item, subunit and equipment unit. The maintainable items have been separated into a number of subunits as shown in Table 5. The proposed structure mirrors one commonly found in asset registers and computerised maintenance management systems in our industry partners. We use the information on where maintainable items sit in the hierarchy to support reasoning in the activity ontology.

4. Reference ontology for maintenance activities

In this section, we present a reference ontology containing of the types of activities that are typically performed in maintenance. This reference ontology contains activity terms collected from 21,088 MWO records and elucidations for the seven core maintenance activities. The methods that we use to extract the activity terms and cluster them according to seven core activities is described in Section 4.1. Elucidations for these activity terms, and the ontological choices underpinning them, are provided in Section 4.2. The purpose of the reference ontology is to provide a set of terms and loose definitions to be used in future ontological analysis by the wider research community. We demonstrate the use of this reference ontology as an input to our application-level model described in Section 5.

Table 4

Example of a data set of maintenance work orders for a single pump over the period of 2012-2020. Abbreviations used: D& E - driver and electrical, PT - power transmission, PU - pump unit, C&M - control and monitoring, LS - lubrication system, P&V - piping and valves.

ID	Date	Unstructured Text	NLP Identified Item (sub-unit)	NLP Identified Activity	Work Order Type	Labour Cost ($)	Material Cost ($)
1	2012-10-18	pressure switch undersize	pressure switch (C&M)	none	corrective	560	821
2	2013-02-14	remove pressure switch	pressure switch (C&M)	remove	corrective	1292	711
3	2013-03-19	calibrate pressure switch	pressure switch (C&M)	calibrate	corrective	413	0
4	2013-04-29	pump not pumping well	pump (PU)	none	corrective	516	0
5	2013-05-21	pressure switch leaking	pressure switch (C&M)	none	corrective	560	700
6	2013-06-03	calibrate pressure switch	pressure switch (C&M)	calibrate	corrective	0	0
7	2013-06-18	18M electrical service motor	motor (D&E)	service	preventative	258	0
8	2013-06-28	calibrate pressure switch	pressure switch (C&M)	calibrate	corrective	322	0
9	2013-08-07	calibrate pressure switch	pressure switch (C&M)	calibrate	corrective	400	0
10	2013-10-06	calibrate pressure switch	pressure switch (C&M)	calibrate	corrective	413	1455
11	2013-10-06	valve needs replaced	valve (P&V)	replace	corrective	650	344
12	2013-12-23	investigate oil leak	oil (LS)	investigate	corrective	0	0
13	2014-02-10	seal leaking	mechanical seal (PU)	none	corrective	3195	5062
14	2014-03-20	pressure switch unserviceable	pressure switch (C&M)	none	corrective	680	1133
15	2014-05-31	check pressure switch	pressure switch (C&M)	check	corrective	145	0
16	2014-07-01	pump not pumping	pump (PU)	none	corrective	64	0
17	2014-11-04	install new pressure switch	pressure switch (C&M)	install	corrective	400	0
18	2014-12-16	78W electrical service motor	motor (D&E)	service	preventative	131	0
19	2014-12-31	26W mech service pump	pump (PU)	service	preventative	0	0
20	2015-03-25	pressure switch failure	pressure switch (C&M)	none	corrective	82	0
21	2015-04-20	replace pressure switch	pressure switch (C&M)	replace	corrective	348	1200
22	2015-05-16	motor tripping on high amps	motor (D&E)	none	corrective	270	0
23	2015-05-19	repair pressure switch	pressure switch (C&M)	repair	corrective	541	0
24	2015-06-05	26W mech service pump	pump (PU)	service	preventative	322	110
25	2015-11-24	26W mech service pump	pump (PU)	service	preventative	478	110
26	2016-06-17	78W electrical service motor	motor (D&E)	service	preventative	91	0
27	2017-06-04	replace pump	pump (PU)	replace	corrective	5168	15370
28	2017-06-12	78W electrical service motor	motor (D&E)	service	preventative	868	0
29	2018-06-15	change oil	oil (LU)	change	corrective	601	0
30	2018-07-03	oil leak from housing seal	mechanical seal (PU)	none	corrective	8632	8369
31	2019-01-06	pump is tripping	pump (PU)	none	corrective	282	0
32	2019-06-06	78W electrical service motor	motor (D&E)	service	preventative	726	0
33	2019-11-26	faulty flowmeter	flowmeter (C&M)	none	corrective	269	0
34	2020-02-19	repair pump	pump (PI)	repair	corrective	3207	566
35	2020-04-14	pressure switch faulty	pressure switch (C&M)	none	corrective	99	0
36	2020-04-22	replace pressure switch	pressure switch (C&M)	replace	corrective	198	2021

4.1. Developing notions for maintenance activity terms

To determine which notions to use for the maintenance activity reference ontology, we follow a two-step process. First, we use a bottom-up approach to examine real-world data and generate clusters of similar activity terms found in this data. Second, we perform a top-down analysis of the clusters by comparing them to existing standards and performing an ontological analysis of the terms.

Table 5

Equipment subdivision for a Centrifugal Pump-Motor System into subunits and maintainable items

Subunit	Driver and Electrical	Power transmission	Pump unit	Control and Monitoring	Lubrication system	Piping and Valves
Maintainable items	Motor Variable drive Power supply	Gearbox Coupling	Pump Casing Impeller Shaft Bearing Mechanical Seal Baseplate Packing	Actuator Controller Sensor Pressure switch Flow switch Flowmeter	Grease Oil Breather Filter	Pipe/Piping Valve Check valve

4.1.1. Step 1: Bottom-up Concept Clustering

To generate a holistic view of the *types* of activities performed by maintenance technicians, we examine 21,088 MWOs, collected over a period of 4.5 years. The work order data looks similar to the sample data shown in Table 1. To extract *activity* terms from these work orders, we use a pre-trained entity recognition model created by the UWA NLP-TLP Group. Extraction of *activity* terms from work order unstructured text is performed using a pre-trained entity recognition model created by the UWA NLP-TLP group [49]. This model is the product of 3 years research effort involving schema development [38], lexical normalization [50], and collaborative annotation [40] of MWOs. The entity recognition model is fine-tuned on a set 6,000 high-quality, gold-standard, annotated MWOs and will extract notions such as *items*, *activities* and *states* from the text. Application of this model to our 21,088 record dataset identified 230 unique root verbs corresponding to maintenance activities after manual lexical normalisation and stemming (e.g. “reeplace" →“replace", and “replacement" →“replace"). These 230 terms are shown in the supplementary materials section in Table 11.

These 230 terms were then clustered according to whether they describe similar activities. For example, the words replace and change out both describe a situation where an item is removed from service and replaced by a new or refurbished item of the same type and make. This was a manual process performed by a subject matter expert. We did attempt word embedding and semantic clustering using automated methods, however due to the close semantics of activity terms, this resulted in an undesirable agglomeration of terms. Each cluster was given a name that best describes the cluster.

We found that some clusters that we developed such as isolate and move do not fit the definition of a maintenance actions described in Section 2.1 (i.e. actions intended to retain an item in, or restore it to, a state in which it can perform a required function” [22]). Therefore, we decided to separate these clusters from *maintenance activities* and instead call them *supporting activities*.

The resulting list of maintenance activity terms and their semantic variants is presented in Table 6 and shown in **bold** font in Table 11 in the supplementary materials. The *supporting activities* are shown unformatted in the same table. While verbs under *supporting activities* constitute a significant portion of our real world data, our focus in this paper is only on the *maintenance activity* words. Future work will consider definitions and competency questions for these supporting activity terms.

4.1.2. Step 2: Top-down Concept Refinement

The engineering community has produced several standards containing lists of maintenance activities, as discussed in Section 2.5. The second step of our process involves refining the initial clusters based on these standards and an ontological analysis. Specifically, we examine ISO14224, and ISO15926-4. The cluster’s name, and membership of similar activities within each cluster, were compared with the terms listed in Table 2. The examination of engineering standards occurred concurrently with our creation of elucidations to define the core maintenance terms tested in this ontology. Refining the clusters required several iterations.

The proposed list of terms is shown in Table 6 and Table 7. These terms diverge from the common terms listed in both ISO14224 and ISO 15926 in the following ways.

Table 6
Maintenance activity terms and semantic variants from the NLP analysis

Maintenance Activity Term	**Semantic variants from the NLP analysis clustered against activity terms**	**Corrective (C) or Preventative (P) activity**
Replace	Change, change out, change over, install, reinstall, re-instate, remove, replace, swap, switch, transition	(C, P)
Repair	fix, free-up, patch, re-engage, reactivate, re-attach, re-build, re-clamp, re-connect, rectify, re-engage, re-fit, refurbish, reinforce, re-line, re-mount, repair, re-route, re-run, re-seal, re-seat, re-secure, reset, re-splice, re-tension, re-torque, re-track, re-weld, re-work, unblock, un-bog	(C)
Inspect	Measure, monitor, read, test, thermography, thickness test	(P)
Adjust	Add, adjust, align, amend, confirm, correct, de-tension, discharge, free-up, guide, position, program, re-adjust, re-align, re-position, review, straighten, tighten, tilt, top-up, track, train	(C)
Service	Charge, clean, drain, fill, grease, hose-out, lubricate, maintain, purge, re-charge, re-fill, rotate, sample, scan, service, tune	(P)
Diagnose	Analyse, assess, check, detect, diagnose, fault-find, find, identify, investigate, re-test, troubleshoot, verify	(C)
Calibrate	Calibrate, re-calibrate	(C,P)

- We have not used the term check. Both check and inspect are described in ISO 14224 as involving investigations characterised by "no maintenance action being performed" [43]. Check is a corrective activity and inspect is a preventative activity so we cannot make the former a subclass of the latter. Instead we propose that check is a synonym for a new class called diagnose. Diagnose is a corrective activity and in addition to capturing check also includes synonyms such as analyse, detect and fault-find.
- We do not use overhaul as a core activity term, instead, we have defined it as a synonym to repair. This is because there is no clear distinction between an overhaul and a repair from the activity alone. Instead, the distinction comes from what is being worked on (i.e. work at the equipment or at the part level). This distinction belongs outside of an ontology designed for describing activities as it depends on several factors that are specific to individual organisations. For example, remote operations may not have the people or equipment to overhaul on site, warranty arrangements might require that overhaul activities be done by the original equipment manufacturer and not by the organisation's maintenance team, and so on.
- We have removed refit and have clustered this word under repair as there are no established rules to distinguish between them other than refit is a small repair.
- We propose that calibrate is a top-level concept as it does not fit into any of our original clusters. Calibration is central to maintenance and involves both the testing and adjustment of instrumentation. Control and instrumentation is a defined subclass in an asset hierarchy [43] and thus can be used in the reasoning. While is is desirable that calibrate is a preventative activity, it can also be corrective.
- We have moved modify to the *supporting activity* terms. In maintenance management modification is considered a on-off activity, and not part of on-going maintenance strategy [51]. Modification work is often done by project groups or contractors external to the regular maintenance teams.

Table 7 provides a subject matter expert definition, a formal definition and an elucidation each of the seven maintenance activity terms. We describe the ontological choices underpinning these elucidations in Section 4.2.

Table 7: Definitions and axioms for maintenance activity terms.

Replace	
SME Def.	Replacement of an item by a new or refurbished item of the same type and make.
Semi-formal Def.	A BFO: Process in which one item is removed and another item with the same required function is installed in its place.
Elucidation	**p** is a *replace activity* = Def. **p** is a *process* and there exist *material artifacts* **a** and **b** such that **a** and **b** participate in **p** at some time. The following is also true: - **a** has a *function* **f** and **b** has a function **f2**. - **f** and **f2** are of type **F** - **a** is a *continuant part of* a *system* **s** at *time* **t** and **b** is not. - **b** is a *continuant part of* a *system* **s** at *time* **t2** and **a** is not. - **t** *precedes* **t2**, **p** *finishes* **t** and **p** *starts* **t2**.
Repair	
SME Def.	Manual maintenance action performed to restore an item to its original appearance or state.
Semi-formal Def.	A BFO:Process, that improves the capability of an item to perform its required function.
Elucidation	**p** is a *repair activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a function **f** and participates in **p** at some time. The following is also true: - **f** is realised in some *functioning process* **fp** - **m** participates at some time in **fp** - **m** has state at some time some *failed state* or *degraded state* **s** and *operating state* **s2** - **s** disables **f**, **p** finishes **s** and **p** starts **s2**
Inspect	
SME Def.	Periodic, careful scrutiny of an item carried out with or without dismantling.
Semi-formal Def.	A BFO:Process, aligned with a preventative maintenance strategy, to estimate the capability of an item to perform its required function.
Elucidation	**p** is an *inspect activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a *function* **f** and participates in **p** at some time. The following is also true: - **f** is realised in some *functioning process* **fp** - **m** participates at some time in **fp** - **m** has state at some time some *state* **s** - **s** is observed in some *observation process* **o** - **p** is prescribed by some part of a *preventative maintenance strategy* **ps**
Diagnose	
SME Def.	The cause of failure is investigated but no maintenance action performed, or action is deferred. Able to regain function by simple actions, e.g.restart or resetting.
Semi-formal Def.	A BFO:Process, performed on observation of an equipment failure, to estimate the capability of an item to perform its required function
Elucidation	**p** is an *diagnose activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a *function* **f** and participates in **p** at some time. The following is also true: - **f** is realised in some *functioning process* **fp** - **m** participates at some time in **fp** - **m** has state at some time some *degraded state* or *failed state state* **s** - **s** is observed in some *observation process* **o**
Adjust	
SME Def.	Bringing any out of tolerance condition into tolerance.
Semi-formal Def.	A BFO:Process involving a change of state in an item that involves no change in capability of the item.
Elucidation	**p** is an *adjust activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a function **f** and participates in **p** at some time. The following is also true: - **f** is realised in some functioning process **fp** - **m** participates at some time in **fp** - **m** has state at some time some *sub-optimal operating state* **s1** and *optimal operating state* **s2** - **p** *finishes* **s1** and **p** *starts* **s2**
Service	
SME Def.	Periodic service task: not normally dismantling of the item.

Semi-formal Def.	A BFO:Process, aligned with a preventative maintenance strategy, involving a change of state in an item that involves no change in the capability of the item.
Elucidation	**p** is a *service activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a function **f** and participates in **p** at some time. The following is also true: - **f** is realised in some functioning process **fp** - **m** participates at some time in **fp** - **m** has state at some time some *operating state* **s** and *operating state* **s2** - **p** finishes **s** and starts **s2** - **p** is is prescribed by some part of a *preventative maintenance strategy* **ps**
Calibrate	
SME Def.	A task involving testing and adjustment of instrumentation to assess and restore function.
Semi-formal Def.	A BFO:Process performed on equipment instrumentation, to estimate the capability of the instrumentation and restore if needed.
Elucidation	**p** is a *calibrate activity* = Def **p** is a *process* and there exists a *material artifact* **m** that has a function **f** and participates in **p** at some time. The following is also true: - **m** is the bearer of a *control and monitoring role* **r** - **f** is realised in some *functioning process* **fp** - **m** participates at some time in **fp** - **m** has state at some time some *failed state* or *degraded state* or *operating state* **s** and *operating state* **s2** - **s** disables **f**, **p** finishes **s** and **p** starts **s2**

4.2. Elucidations for seven core activity terms

In this section, we discuss the elucidations given in Table 7 and the ontological choices underpinning these elucidations. First, we describe the eludication for the replace activity. This elucidation describes a situation where there are two material artefacts (**a** and **b**) that have functions of the same type, **f**. When the replace activity occurs, **a** is a continuant part of a system before the replacement, and **b** is a continuant part of the system after the replacement. This is true whether a whole piece of equipment is replaced (i.e. a pump) or only part of a system is replaced (i.e. a pressure switch). A critical part of this definition is that both the item to be replaced (**a**) and the item to be installed (**b**) both fulfil the same *function* when installed as part of a system. In other emerging standards such as ISO15926 Part 14 [52], there is the concept of a *functional object* and the relation *functionalPartOf*. In BFO, there is no such concept and this makes it difficult to distinguish where the replaced asset sits in a functional breakdown (as opposed to a physical breakdown). For our purposes, we believe it is sufficient to say that the function of the asset is the same for both **a** and **b**. In creating this elucidation, we made a choice about whether a material artefact can have a function **f** if it is not yet installed in the system. For example, if **b** is soon to be installed, does it have **f** as a function before the replace activity begins? We argue that the function **f** is still the designed purpose of the item, regardless of whether it is installed or not. Instead, the “installation status" (or the state of the item) will impact whether or not that function can be *realised*, as per our previous work on maintenance state [53]. Finally, we acknowledge that the idea of replacement is a deeply philosophical problem and raises many questions about identity [54]. Therefore, this elucidation does not make any assertions about the identity of an asset, before and after components have been replaced.

The semi-formal definition for a repair activity describes a process that improves the “capability of an item to perform its required function". For this, and for the maintenance terms that follow, we draw on our previous work on the notion of maintenance state [53]. In this work, we model the states in which items do not have the ability to perform functions (i.e. participate in a process that realises a specific function). We achieve this through a *disables* object property where some *state disables* a *function realising process*. Furthermore, we determine that in maintenance, when an item is installed, but does not have the capability to perform its required function, then it is in a *degraded state* or a *failed state*. Otherwise, the item is in an *operating state* and can perform all of its functions. Given this work, we describe a repair activity as a situation where there is a material artefact **m** that participates in this activity. **m** has some function **f** that is realised in some functioning process **fp**. Before the repair activity, **m** is in some degraded state or failed state **s** (i.e. a state that disables **f**). After the repair activity, **m** is in a *operating*

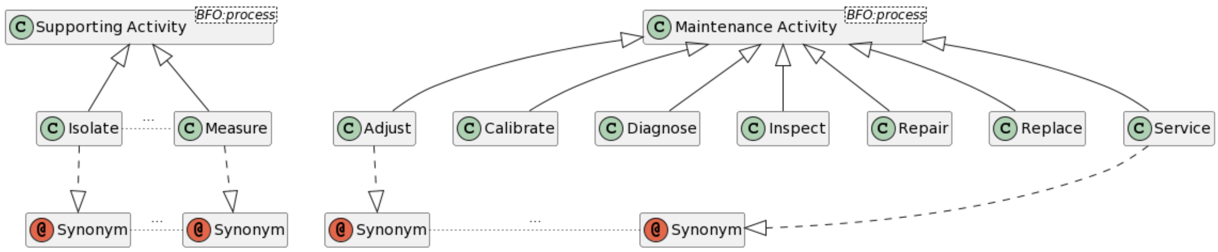

Fig. 1. A conceptual diagram of the maintenance activity reference ontology (synonyms are listed in Table 6)

state. Using our modelling of state, we have been able to model changes of state and their implications a result of a replace activity.

An inspect activity is described as a process to *estimate* the capability of an item to perform its required function. Since an item's *state* determines its ability to perform its required function, as described previously, this elucidation describes an *observation process* **o** in which the item's *state* **s** is examined. Note that from a maintenance perspective, an item's ability to perform its required function is not something that is measures (or estimated) on a continuous scale. Instead, the item is either able to perform its function or it is not. Therefore, we believe an observation of the item's state (failed, degraded or operating) is sufficient to model an *inspect activity*.

An inspect activity differs from a diagnose activity in one crucial aspect. In the SME definition, an inspect activity is defined as "periodic". This means, that a preventative maintenance strategy is defined that determines at what intervals an item should be inspected. A diagnosis, on the other hand, is a *corrective* action to determine "the cause" of a failure (and is not scheduled). Note that even though an action is *corrective*, it may still be described in a maintenance strategy in the case of a Run-to-Failure corrective maintenance strategy. This distinction between corrective and preventative activities is crucial to our application-level ontology described in Section 5. Without these elucidations and the analysis performed in this section, it is very difficult to distinguish between a diagnose and inspect activity.

An adjust activity is a process that involves changing the state of an item but with no change to the capability of the item. This means that adjust activities are performed before an item would enter a degraded state. The elucidation for an adjust activity looks very similar to a repair activity as they are both corrective actions involving a change of state. The difference is that a *repair activity* requires the item to be in a *failed state* or *degraded state* before the activity (i.e. it is unable to perform a required function). Instead, an adjust activity will remain in an *operating state* throughout **p**. This means that there is no change to the item's ability to perform functions before or after the activity. Whilst in this *operating state*, the *material artefact* **m** will go "out of tolerance", as per the SME definition in Table 7. Therefore, we introduce two sub-classes of *operating state* (extending our work in [53]). We model **s1** as a *sub-optimal operating state* (i.e. where equipment is out of tolerance) and **s2** as an *optimal operating state*. As with an inspect activity and a diagnose activity, the difference between an adjust activity and a service activity is that an adjustment is a corrective action, and a service occurs at periodic intervals and is aligned with a preventative maintenance strategy.

The final maintenance activity type is a calibration activity. This is defined as a task involving testing and adjustment of instrumentation to assess and restore its function. In our elucidation, we define a calibration activity similarly to a repair activity, but the item can be in a functioning, degraded or failed state before the activity. What distinguishes a calibration from other activity types is that it is performed on *instrumentation*. Therefore, in the elucidation, we say that the material artefact **m** plays a control and monitoring role **r**. This relationship is important for our use case in Section 5, as we distinguish calibrate activities by their membership in a Control and Motoring subunit of the centrifugal pump. Finally, a calibrate activity can be either corrective or preventative so we have made no assumptions about whether or not it is is prescribed by some part of a preventative maintenance strategy (as with service and inspect activities).

4.3. Reference ontology Implementation

The reference ontology is implemented in OWL and is available at https://github.com/uwasystemhealth/Paper_Archive_Maintenance_Activity in the file *maintenance-activity.owl*. A conceptual diagram of this reference ontology is shown in Figure 1 (where "C" is a class and "@" is an annotation property). In the ontology, there are two classes, Maintenance Activity and Supporting Activity. These activities are a subclass of BFO:Process. The core maintenance activities and supporting activities that were identified in Section 4.1.1 are classes in the ontology and have their SME definition, semi-formal definition and elucidation stored as annotation properties. The other terms that formed part of each cluster (i.e. align and amend as members of the adjust cluster) are stored in an annotation property, synonym. "Synonym" is defined in the IOF as "an alternative label (designation) used for the resource in some community" [55]. While it could be argued that the words listed as synonyms are not "true" synonyms with one another in an English language sense, the maintenance community interprets these words in the same way. Therefore, for our maintenance activity reference ontology, this is an appropriate annotation property to use.

5. Application-level ontology for MWO data quality

In this section we define an application-level ontology to answer a set of competency questions related to the use case described in Section 5.1. The purpose of this model is to meet a real industrial need while making use of the reference-level terms and ontological analysis described in Section 4. As discussed, the quality of maintenance work order data is an unsolved problem that is imperative for downstream statistical tasks such as mean-time-to-failure analyses. However, maintenance engineers rarely spend time checking the quality of this data, as there is no widely-accepted, repeatable manner to perform this task. This model uses SWRL rules and ontological classification to provide a structured and repeatable way for organisations to assess the quality of maintenance work orders that reflects the logic currently used by engineers performing this task in industry.

5.1. Competency questions

The application-level ontology is designed to perform a specific reasoning task to meet a real industrial need. The work order data described in our use case (Section 3), contains eight distinct activity words. These are remove, calibrate, service, replace, investigate, check, install, and change. For data quality purposes, we want to determine if the activity term is consistent with other information in the MWO. The competency questions that we use to assess this ontology are:

Q1. Check classification of records with the activity replace.

Q2. Check classification of records with the activity calibrate.

Q3. Check classification of records with the activity investigate.

Q4. Check classification of records with the activity service.

The logic used by engineers when performing these data quality checks is shown in Figure 2. First the item that work was performed on is identified. This might be at the equipment level e.g. pump, or at the maintainable item level e.g. pressure switch. Items at the maintainable item level can be located in their respective sub-classes using standardised asset class hierarchies as shown in Table 5. Next a check is made to see if the labour cost field has a value greater than zero dollars. This is necessary to determine if the proposed activity occurred or not. Reliability engineers are interested in both cases. Our interest here is only on MWO's that are actioned. The next decision point is determining if the *work type* is preventative or corrective. This allows for a split into two groups. Repair, adjust and diagnose are activities arising from corrective strategies. Inspect and service are activities arising from preventative strategies. Replace and calibrate can be either. The next stages of the logic in both branches depends on

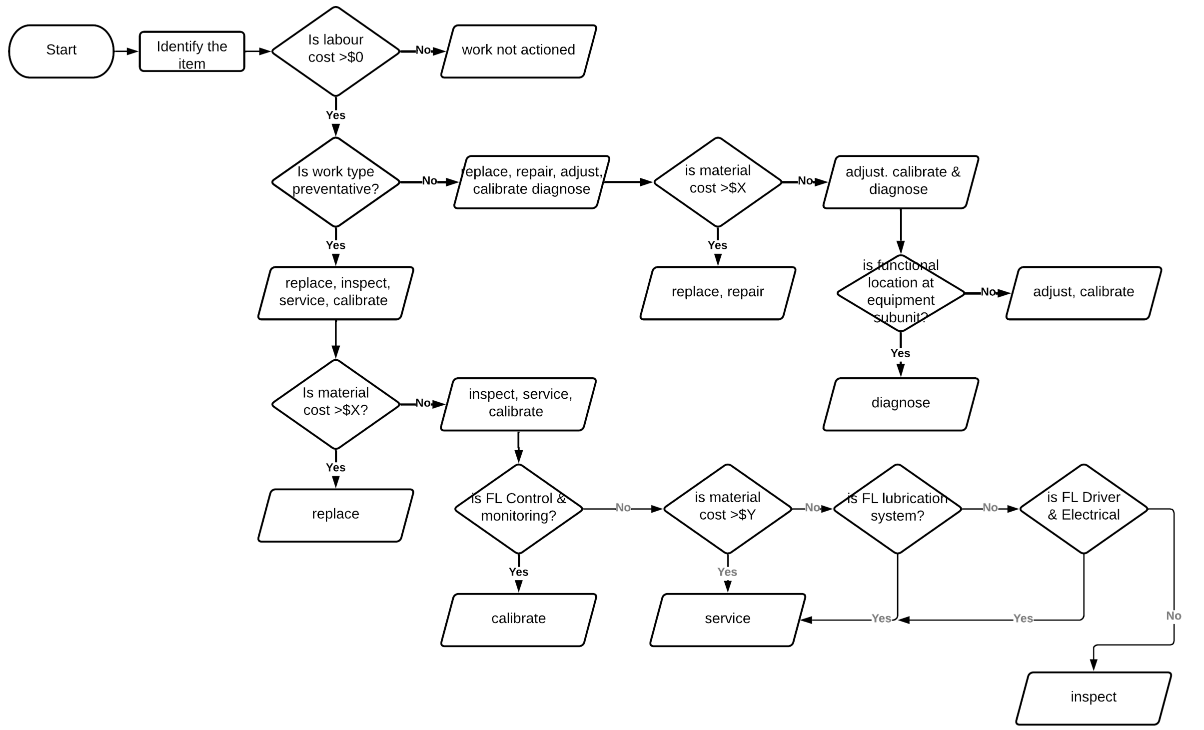

Fig. 2. Decision tree capturing the logic of a typical reliability engineer checking the validity of maintenance activity terms in MWOs. The X and Y values for the \$ amount depend on the item; $X >> Y$

information in the *material cost* field and on the *item*. For instance, if material costs are incurred then the activity is not an inspection. Distinguishing between the classes of service and inspection is done based on the item and which subunit it belongs to. Calibrate is, for example, associated with items in the Control & Monitoring subclass. The reasoning here is representative of steps taken by reliability engineers every day as they manually process MWOs [15].

5.2. *Application-level Ontology Implementation*

This ontology was developed making use of the reference ontology described in Section 4. The ontology consists of 8 modules, as shown in Figure 3. A description of each of these modules is contained in in Table 8. This ontology can also be found on GitHub at https://github.com/uwasystemhealth/Paper_Archive_Maintenance_Activity. In the following sections, we will discuss the key ontological choices made in the development of this ontology.

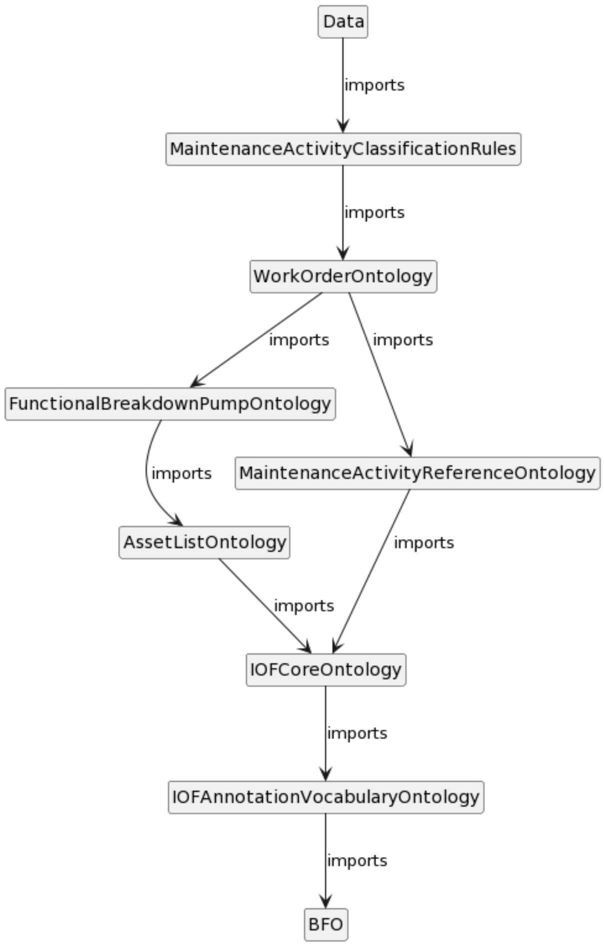

Fig. 3. Import Structure of Application-level ontology for assessing MWO data quality

Table 8
Purpose of each module in the Application-level ontology for assessing MWO data quality

Module	**Description**
Data	Contains individuals extracted from MWO data in Table 4.
Maintenance Activity Classification Rules	Contains SWRL rules, and supporting classifiers, used to perform activity classifications according to the the decision tree in Figure 2.
Work Order Ontology	Contains notions related to MWOs (i.e. Maintenance Work Order Record, Maintenance Work Order Record Data Item, Work Order Execution Process).
Functional Breakdown Pump Ontology	Imports the assets contained in the Asset List Ontology and divides them into into the subunits shown in Table 5.
Asset List Ontology	Contains a list of material artefact classes in a centrifugal pump (e.g. Mechanical Seal).
Maintenance Activity Reference Ontology	The reference ontology for maintenance activities described in Section 4.
IOF Core Ontology	The IOF Core Reference Ontology retrieved from [55].
IOF Annotation Vocabulary Ontology	The IOF Annotation Vocabulary (includes Synonym, for example) Ontology retrieved from [55].
BFO	The Basic Formal Ontology, used as a foundational ontology [27].

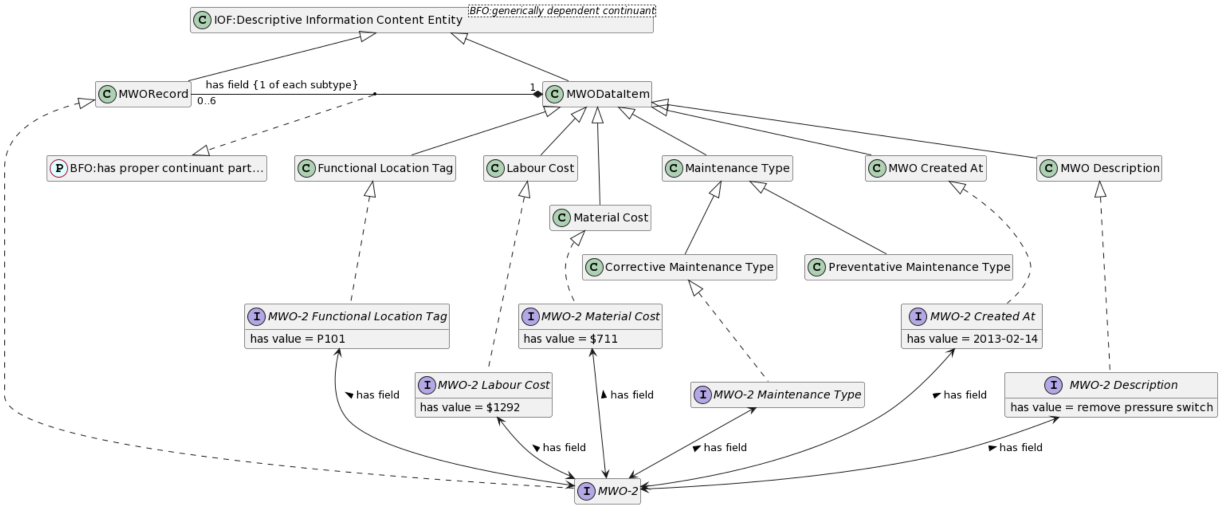

Fig. 4. A conceptual diagram of information content entities for work order records in the application-level ontology

5.2.1. *Work Order Records and Information Content*

At the core of this application-level ontology is an Information Content Entity, the MWO Record. In this section, we will explain our choices in modelling a MWO Record and its information parts (i.e. Material Cost, Labor Cost). A conceptual diagram of the information content entities in the application-level ontology can be found in Figure 4 (where “C" denotes a Class, “I" denotes an Individual and “P" denotes a property). Note that we are choosing not to use the Information Artefact Ontology as a reference ontology, as it does not contribute to the ontology for this use case.

As shown in Table 4, a MWO Record has multiple fields (or MWO Data Items that contain information relevant to the MWO. To capture the information contained in these fields, we have added a new object property called refers to that is a sub-property of is about. If an Information Content Entity refers to an entity, then it uses, evokes, or references the entity or concept in some way that is not necessarily descriptive, designative, nor prescriptive.

In this ontology, refers to is used in several cases including:

- Where a Descriptive Information Content Entity, e.g., a Work Order Record field, references an asset in a functional location; the functional location tag field of the record cannot be a designative entity in this case as the designation of the functional location (and/or the asset installed in the location) comes from elsewhere, the field merely contains a value that matches the designation of the location/asset and, hence, refers to it. This could be a Referential Information Content Entity. Moreover, the field cannot be said to describe the functional location/asset as it is describing an aspect of the work order process/activity, by identifying a primary participant in the process/activity.
- Where a Descriptive Information Content Entity such as a Work Order field, is deduced to be referencing an entity or concept, e.g., the work order record description field having been identified (through the NLP pre-processing) as referring to a specific *activity* type, which is reconciled from the textual form used in the description and the terms associated with the different activity classes.

Finally, we are choosing not to link the MWO Record (as a generically dependent continuant) entities to their concretizations in the originating database. Instead, we provide an abstraction /normalisation over the records stored in an actual database. We take for granted, at the moment, that there is an originating database and, hence, some specifically dependent continuant (i.e., the physical datum) and the common independent continuant to which both the datum and the MWO record relate.

5.2.2. *NLP Identified Activities*

As discussed, the goal of this application-level ontology is to check if the activity described in the MWO's unstructured text field matches the information in the rest of the ontology. To achieve this, we store the NLP Identified

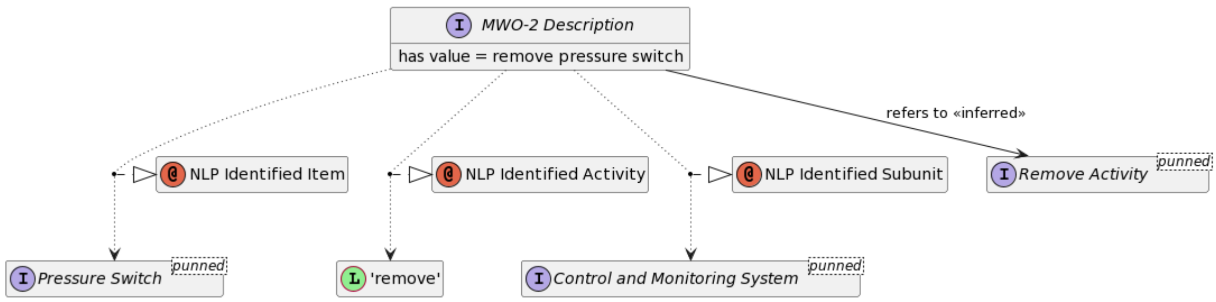

Fig. 5. A conceptual diagram of NLP identified properties for work order descriptions in the application-level ontology

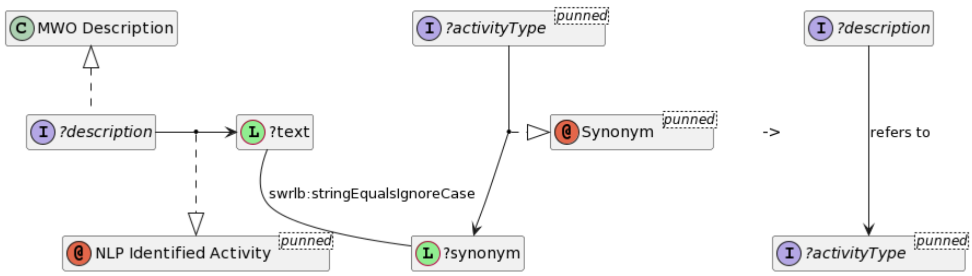

Fig. 6. SWRL rule reconciling NLP extracted activity term with the ontology. Elements representing variables have names beginning with a question mark (‘?’).

Activity from Table 4 as an annotation property on the Maintenance Work Order Description. A conceptual diagram of this pattern in shown in Figure 5 (where “L" represents a ‘Literal’ property value). We represent this information in annotation properties because they represent *assumed* knowledge resulting from an entity recognition algorithm applied to the specific field. Thus we cannot assert this knowledge as individuals in the ontology without a detailed ontological analysis, which is not necessary for our use case. We use the information stored in this annotation property to infer the NLP identified activity type in the maintenance activity reference ontology described in Section 4. We achieve this inference using punning of the property and activity types, and a SWRL rule that compares the normalised text value of the annotation property against the synonyms of the activity classes. We reconcile term and activity type within the ontology as the term clusters and annotations may be updated in the ontology based on an organisation’s profiling of their own data. The SWRL rule is illustrated in Figure 6.

Using the inferred maintenance activity, we can then compare the NLP identified activity to the result of the decision tree in Figure 2. The rules for which are discussed in Section 5.3. If there is no NLP Identified Activity, the MWO record is considered to refer to an Unspecified Activity. We acknowledge that MWO records with an Unspecified Activity lead to records being considered ‘non-matching’ in the comparison; however, beneficial analysis can still result by flagging such records for review by an engineer. In the following section, we will describe how the NLP Identified Item and the NLP Identified subunit are used to gain further information about the *true* activity performed in the work order.

5.2.3. *Items and subunits*

To realise the logic contained in the decision tree in Figure 2, we need to know the type of equipment being worked on (or, more specifically, which subunit from Table 5 the item fits into). In the MWO structured data, there is a functional location field (a Functional Location Tag in our ontology). For all of the work orders in this use case, the Functional Location Tag refers to a Pump System (a centrifugal pump and other equipment e.g, motor, values, instrumentation necessary to perform the pumping function). The Pump System and other item classes are defined as IOF Core Material Artifacts, which is a BFO Object class. In our ontology, a Pump System has its (functional) subunits (e.g. the Control and Monitoring System) as continuant parts *at all times*. The subunits, modelled as Engineered System classes under the BFO obect aggregate hierarchy, have continuant parts *at some*

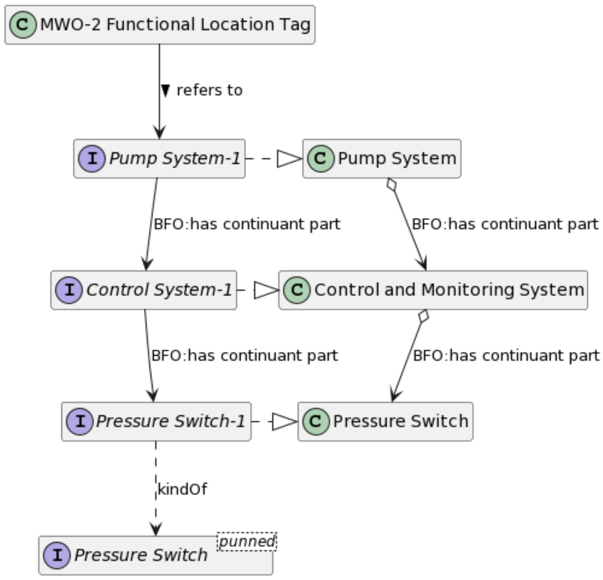

Fig. 7. A conceptual diagram of items and subunits in the application-level ontology

time, e.g., the Control and Monitoring System having further continuant parts such as a Pressure Switch and other types of instrumentation. This is a minimal representation of the asset's functional breakdown structure, which we do not address further as it is not the focus of this paper and is a complex topic to address in BFO. A conceptual view of items and subunits in the application-level ontology is shown in Figure 7. For reasoning purposes, the item and subunit classes are punned to individuals and matching values for the object property *kindOf* are inferred by the classification. The equipment items are parts, the classes for which are contained in the Asset List Ontology module while the equipment's classes membership in a subunit is captured in the Functional Breakdown Ontology (as described in Table 8)—the punning and *kindOf* property are separated into the Maintenance Activity Classification Rules module.

The only structured information in the MWO record that indicates the item being acted upon is the Functional Location Tag. Since this is not at the right level of granularity to perform our analysis, we rely on our entity recognition algorithm, and some subsequent analysis, to identify the specific item and subunit. For the purpose of this analysis, we use a simple mapping from the item type to subunit type. The breakdown structure of centrifugal pumps identified in Table 5 has a unique subunit type for each maintainable item type. Using the NLP Indentified Item (model shown in Figure 5), we can utilise the subunit information in the rules to identify possible participants of the activity, i.e., the item individuals at the part level, and further use this to infer a classification for the maintenance activity that was performed on the item.

5.3. *Maintenance Work Order Execution Processes and Activities*

Thus far, we have shown how we model the structured information from the work orders as Information Content Entities. Subsequently, we have shown how we infer the *expected* maintenance activity that was performed on the item (from the NLP Identified Activity). Finally, we have shown how we use extra data from our unstructured text by storing the item being acted on and the item's subunit. In this section, we show how we use this information to realise the logic shown in Figure 2 and infer the type of maintenance activity that was likely to have been performed. If the type of maintenance activity identified through analysis of the description does not match the inferred maintenance activity classification, we assume that the MWO data is not of a high quality and requires further human attention.

The conceptual diagram in Figure 8 shows how we model the inference of the activity type and its relation to a MWO record. In the figure, our Maintenance Work Order Record (MWO-2) describes some Maintenance Work Order Execution Process, following the structure of the ROMAIN ontology [30]. This execution process represents

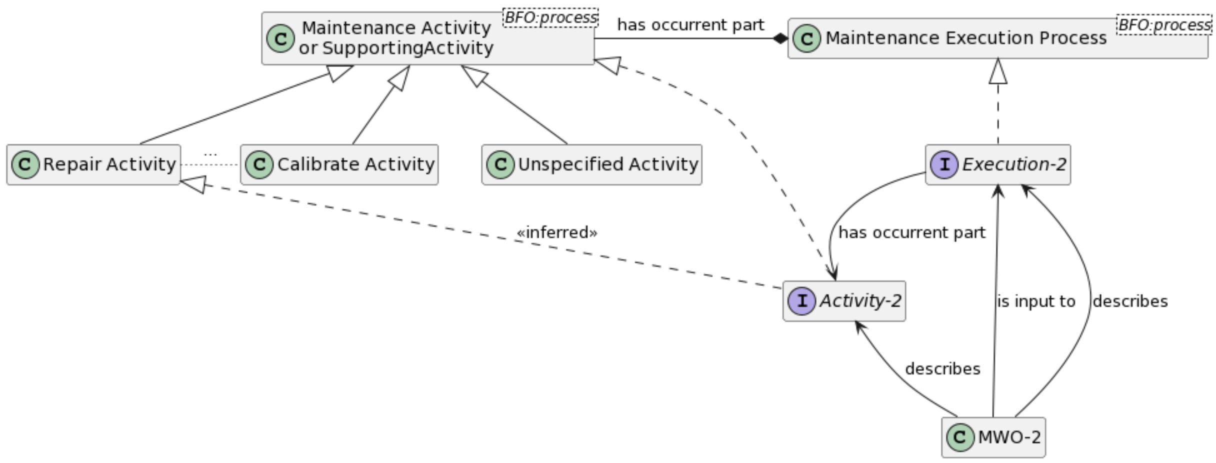

Fig. 8. A conceptual diagram of maintenance work order execution processes and activities in the application-level ontology

all the activities that are carried out to execute a MWO (and is contained in the Maintenance Work Order Ontology from Table 8). This execution process has the Maintenance Activity to be classified as an occurrent part. In the case of an activity that was not actioned, the classification is applied to the execution process and not the activity. This is because the activity cannot really exist if the activity to be classified was not actioned, but other parts of the process may have been. This way the individual for the unactioned activity could be removed if desired.

The actual maintenance activity classification is inferred using SWRL rules that capture the logic of the decision tree shown in Figure 2. To support this inference, there are several supporting elements, including additional classifiers and rules, that surface information used by other rules or for subsequent querying. The primary rules and supporting elements are captured in the Maintenance Activity Classification Rules module, described in Table 8.

5.3.1. Additional Activity Classifiers

There are two main categories of additional activity classifiers: intermediate classifiers, and final classifiers. The latter correspond to classifications inferred at the leaf nodes of the decision tree (refer Figure 2), while the former may be inferred at non-leaf nodes. Furthermore, some of the final classifiers are intended for querying purposes while others represent the outcome of the decision logic.

The interim classifiers are based on the subsets of activity class they encompass and include:

- Maintenance Supporting or Unspecified Activity (Maintenance Activity, Supporting Activity, or Unspecified Activity)
- Corrective Action (adjust, calibrate, diagnose, repair, or replace)
- Corrective Action with Low Material Cost (adjust, calibrate, or diagnose)
- Preventative Action (calibrate, inspect, replace, or service)
- Preventative Action with Low Material Cost (calibrate, inspect, or service)

While the Corrective Action and Preventative Action classifications in the rules could have been inferred through an OWL *equivalent class* axioms, it was chosen to use rules for consistency with the other rules that could not do so.

The final classifiers include a pair of disjunctive classifiers, to meet the decision logic, as well as maker classes that indicated the status of the classification. These include:

- Adjust or Calibrate—disjunctive class indicating that the activity may be either an adjust activity or a calibrate activity, but the currently available information and decision logic cannot ascertain for certain
- Repair or Replace—disjunctinve class indicating that the activity may be either a repair activity or a replace activity, but the currently available information and decision logic cannot ascertain for certain
- Inferred Activity—classifier marking activities that have had their classification inferred; is separated into several subclasses:

* Inferred Type Matches Work Order Description—classifier marking activities as matching the NLP Identified Activity of the MWO record's description field
* Inferred Type Does Not Match Work Order Description—classifier for marking activities as not matching the MWO's description field; due to the restrictions of OWL DL and SWRL, this is not used by the rules but can be used by subsequent queries and assertions to track such activities more readily
* Uncertain Maintenance Activity—classifier marking activities that are not clearly distinguished as a specific activity type, for example, it will accompany Adjust or Calibrate and Unspecified Activity classifications

5.3.2. Additional Equipment Classifiers

For ease of writing the SWRL rules, some additional classifiers are used to identify particular subsets of equipment subunits required by the rules. The classes are quite general and apply orthogonal classifications without restricting the primary classification of subunit. As such, they should be broadly applicable to categories of equipment or, at the very least, the classes of pump other than centrifugal pump. This should control the number of classes required when defining activity classification rules for broader classes of equipment than illustrated here. For example, it does not matter if a particular class of pump does not include a subunit of a particular type; only if there is a contradiction in the types of subunit relevant to particular actions would more fine-grained distinctions need to be made. While it is possible to directly identify the class expressions in the rules, providing named classes is more flexible, convenient, and manageable. The classes include:

- Pump Subunit—disjunction of Engineered System subclasses that are can be included in Pumps as subunits.
  * Inspectable Unit—subunit classifier identifier subunits that are relevant to inspection activities, i.e., not Lubrication System, nor Control and Monitoring System, nor Driver and Electrical System
  * Not Control And Monitoring System—some rules make reference to any subunit other than a Control and Monitoring System: this class captures that requirement
  * Not Pump Unit System—some rules make reference to any subunit other than the Pump Unit System: this class captures that requirement

5.3.3. Item Material and Service Costs

The decision logic for classifying the activities relies quite heavily on the material costs of the work that was performed. According to the decision tree of Figure 2, there are two different material costs that are evaluated X and Y, called here *material item cost* and *item service cost*, respectively. The material item cost is an indicator of the expected minimum cost if an item were to be replaced/repaired, while the service cost is an indicator of the minimum costs associated with servicing an item. A pair of data properties is used to capture this information on a per item basis (likely derived from the item's class). Such an approach makes the rules more flexible as specfic values are not hardcoded into them.

Of course, the item cost and service cost information is not currently in the scope of the available data. However, this information could be derived from other sources, data profiling, or other means of estimating the costs associated with activities related to equipment. If available, such information can be incorporated without modifying the rules to improve the granularity of the activity classification rules. Since we do not currently have the detailed information, we assume that the current data is indicative of the maintenance activities performed on the centrifugal pumps of this organisation and set default values accordingly: material item cost is \$150, while item service cost is set at \$1.

The possibility of avoiding the material costs, as it pertains to replacements, is discussed in Section 7.1.2 as it is unlikely to be the most accurate approach in general.

5.3.4. SWRL Rules

In total there are 13 rules for classifying activities according to the leaf nodes of the decision tree (Figure 2), 4 additional rules inferring intermediate classifications, and 2 rules for final comparison of classification to the NLP Identified Activity type. Figure 9 illustrates a simple rule for identifying the corrective maintenance branch of the decision tree.

Intermediate classifications cannot be fully relied upon due to some of the activity classes falling under both corrective and preventative maintenance categories. This application ontology does not address the ontological concern of the activity individual being a corrective or preventative occurrence, it merely uses the classifications

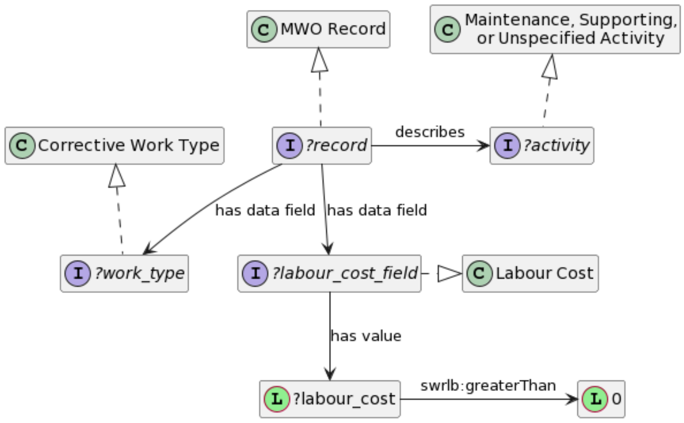

Fig. 9. Rule for inferring the Corrective Action classification based on the information in the MWO record

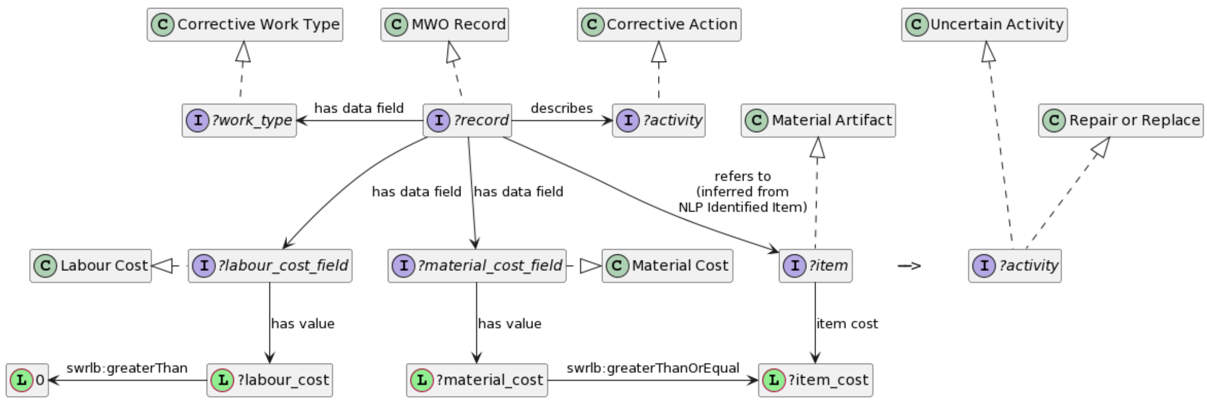

Fig. 10. Rule inferring the Repair or Replace classification, including the marker classification Uncertain Activity, based on the information in the MWO record.

as indicators. Therefore, the actual maintenance type of the associated MWO record is relied upon to ensure the final classifications on the correct path of the tree where activity classes that may be of either category are involved. Capturing it in a rule also makes it easier to reuse the pattern in other, more complex, rules.

The rule for the final cross-check for matching the activity classification is shown in Figure 11.

Querying the classifications with SPARQL Using the inferred information the inferred activity classifications can be queried including but, not limited to, all inferred activity classifications, activity classifications matching the NLP Identified Activity information, uncertain activity classifications, and non-matching activity classifications. Of particular interest are the non-matching, including uncertain, activity classifications. Due to the constraints of OWL DL and SWRL, retrieving the non-matching activity classifications leverages the Inferred Type Matches Work Order Description class and SPARQL negation to match the non-existence of the pattern in which it occurs. An example query is shown in Figure 12.

The query identifies all the activity individuals that have an inferred activity classification, a Maintenance Activity class, that does not match that class indicated by the NLP Identified Activity. As output, it includes the classifications of the activity individual as well as the activity class associated with the MWO record (via the NLP Identified Activity property). The latter part of the query is optional so as to include results where the activity class associated with the MWO record is missing, in case there is an error in the process. Such a query allows an engineer to review the non-matching, including uncertain, classifications against what was written in the MWO record; they can then take action to rectify errors to improve the data quality.

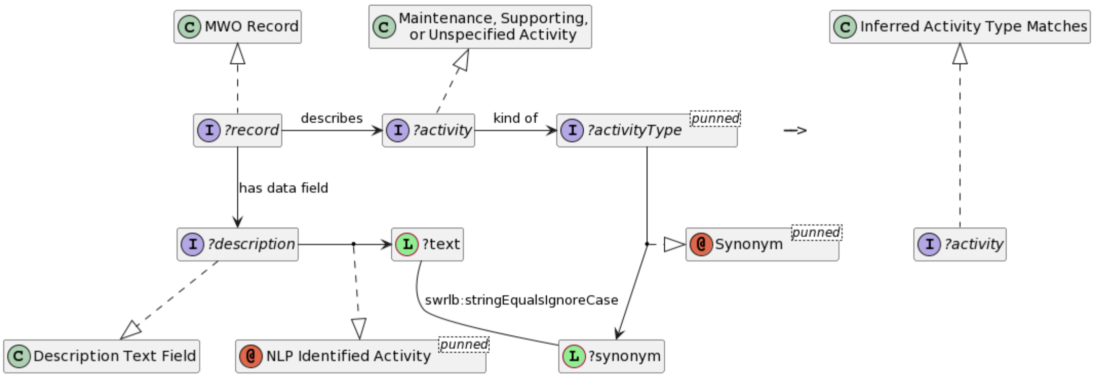

Fig. 11. Rule inferring the match between a concrete activity classification and the NLP Identified Activity type

Fig. 12. SPARQL Query for retrieving the activities that did *not* match the NLP Identified Activity type

```
PREFIX core:<http://www.industrialontologies.org/core/>
PREFIX activity:<http://www.semanticweb.org/maintenance-activity#>
PREFIX rule:<http://www.semanticweb.org/maintenance-activity-classification-rules#>
SELECT DISTINCT ?activity ?clazz ?nlpActivityType
WHERE {
  ?activity a/rdfs:subClassOf* rule:InferredActivity ;
    a ?clazz .
  FILTER ( ?clazz != rule:InferredActivity &&
           ?clazz != rule:UncertainActivity &&
           ?clazz != rule:MaintenanceSupportingOrUnspecifiedActivity )
  FILTER NOT EXISTS {
    ?activity a/rdfs:subClassOf* rule:InferredActivityTypeMatchesWorkOrderDescription
  }
  OPTIONAL {
    ?activity core:describedBy|^core:describes ?_record .
    ?_record rule:refers_to ?nlpActivityType .
    ?nlpActivityType rdfs:subClassOf+ activity:MaintenanceActivity .
  }
}
```

6. Evaluation

We evaluate this ontology by assessing its performance against the competency questions defined in Section 5.1. The answers produced the by the ontology reasoning processes were compared with the decision tree used by the subject matter expert shown in Figure 2. Each competency question is addressed in turn in the text section below. In addition, Table 9 is used to identify data quality issues in each record. The table shows for each record a) the activity term in the original MWO unstructured text, b) the activity type resolved from the semantic synonyms captured in the maintenance reference ontology, and c) the activity classification inferred by the ontology and rules. There are two types of data quality issues to identify in this analysis. The first is missing data, in this case where the creator of the MWO record failed to include an activity word in their description. For example in the MWO with ID 5, the MWO Description is ‘pressure switch leaking’. The second is inconsistency between the identified maintenance activity cluster term and the inferred activity classification. For example in the MWO with ID 10, ‘calibrate pressure switch’, the resolved activity from NLP is a calibrate activity but the ontology suggests this is a Repair or Replace activity as there are material costs of $1455. Material costs of this magnitude for a pressure switch are inconsistent with only a calibration being performed (according to the logic given in Figure 2). We suggest that

Table 9
Results of activity classification by the ontology and rules

ID	Activity Term from the original record	Clustered Activity Type (from NLP Identified Activity)	Activity Classification
1	none	—	Repair or Replace
2	remove	replace	Repair or Replace
3	calibrate	calibrate	Adjust or Calibrate
4	none	—	diagnose
5	none	—	Repair or Replace
6	calibrate	calibrate	Work Not Actioned
7	service	calibrate	calibrate
8	calibrate	calibrate	Adjust or Calibrate
9	calibrate	calibrate	Adjust or Calibrate
10	calibrate	calibrate	Repair or Replace
11	replace	replace	Repair or Replace
12	investigate	diagnose	Work Not Actioned
13	none	—	Repair or Replace
14	none	—	Repair or Replace
15	check	diagnose	Adjust or Calibrate
16	none	—	diagnose
17	install	replace	Adjust or Calibrate
18	service	service	service
19	service	service	Work Not Actioned
20	none	—	Adjust or Calibrate
21	replace	replace	Repair or Replace
22	none	—	Adjust or Calibrate
23	repair	repair	Adjust or Calibrate
24	service	service	service
25	service	service	service
26	service	service	service
27	replace	replace	Repair or Replace
28	service	service	service
29	change	replace	Adjust or Calibrate
30	none	—	replace
31	none	—	Adjust or Calibrate
32	service	service	service
33	none	—	Adjust or Calibrate
34	repair	repair	Repair or Replace
35	none	—	Adjust or Calibrate
36	replace	replace	Repair or Replace

once the maintenance technicians started the job that they realised the pressure switch needed to be replaced, and did so. There is no facility for technicians to re-word the original MWO text hence the need for these data quality checks.

Question 1: Check classification of records with the activity replace.

Answer using decision tree:

- MWOs with IDs 2, 11, 21, 27 and 36 contain the replace, or a semantic synonym (e.g remove), in their unstructured text. Examples include ID 2 ‘remove pressure switch’ and ID 11 ‘valve needs replaced’. All are corrective actions with material costs, so should be classified as Repair or Replace.

- MWO with ID 29 is a corrective action involving a change of oil. There is no material cost, hence, it is likely that the oil was not actually changed and was instead an adjust or calibrate activity that was performed.

Answer using ontology:

- MWOs with IDs 2, 11, 21, 27 and 36 have all been classified as Repair or Replace. All NLP Identified Activities have been resolved as the activity replace. Therefore, there are no data quality issues identified by the ontology for these records.
- MWO with ID 29 has been classified as Adjust or Calibrate. The NLP Identified Activity has been resolved as replace. Due to a mismatch between the classified activity and the resolved activity, a data quality issue has been identified by the ontology. This record should be reviewed by an engineer.

Question 2: Check classification of records with the activity calibrate.

Answer using decision tree:

- MWOs with IDs 3, 8 and 9 have the activity calibrate in the unstructured text. These are corrective actions with no material costs. The *item* (pressure switch) is at the maintainable item level so the logic classifies these MWOs as Adjust or Calibrate.
- MWO with ID 6 has the activity calibrate in its unstructured text but zero labour cost, hence the classification should be Work Not Actioned.
- MWO with ID 10 has the activity calibrate in its unstructured text. It is a corrective action with a material cost of $700. The logic classification for this is a Replace or Repair.

Answer using ontology:

- MWOs with IDs 3, 8 and 9 have all been classified as Adjust or Calibrate. All NLP Identified Activities have been resolved as the activity type calibrate. Therefore, there are no data quality issues identified by the ontology for these records.
- MWO with ID 6 has been classified as Work Not Actioned. The NLP Identified Activity has been resolved as calibrate. Due to a mismatch between the classified activity and the resolved activity, a data quality issue has been identified by the ontology.
- MWO with ID 10 has been classified as Repair or Replace. The activity type has been resolved as calibrate. Due to a mismatch between the classified activity and the resolved activity, a data quality issue has been identified by the ontology.

Question 3: Check classification of records with the activity investigate.

Answer using decision tree:

- MWO with ID 12 contains the activity investigate in its unstructured text. However, as this contains no labour cost the logic classifies this as ‘work not actioned’.

Answer using ontology

- MWO with ID 12 has been classified as Work Not Actioned. The NLP Identified Activity has been resolved as the activity type diagnose. Due to a mismatch between the classified activity and the resolved activity, a data quality issue has been identified by the ontology.

Question 4: Check classification of records with the activity service.

Answer using decision tree:

- MWOs with IDs 18, 26, 28 and 32 are preventative actions involving a 78W (every 78 weeks) electrical motor service. They all have zero material cost consistent with an electrical service. The motor is part of the Driver and Electrical subunit and so the logic classifies these activities as service.

Table 10
Data quality performance

MWO records	Records with missing activity term in the unstructured text but an inferred activity term provided by the ontology	Records with inconsistencies in the activity term	Records with no data quality issues
IDs	1, 4, 5, 13, 14, 16, 20, 22, 30, 31, 33, 35	6, 10, 12, 15, 17, 19, 23, 29	2, 3, 7, 8, 9, 11, 18, 21, 24, 25, 26, 27, 28, 32, 34, 36
Number (n)	n = 12	n = 8	n = 16

- MWOs with IDs 24 and 25 are preventative actions involving 24W mechanical service pump. Each has an associated material cost of less than \$150. These are therefore classified in the logic as the activity service.

Answer using ontology:

- MWOs with IDs 18, 24, 25, 26, 28 and 32 have all been classified as a service activity. All NLP Identified Activities have been resolved as the activity type service. Therefore, there are no data data quality issues identified by the ontology for these records.

7. Discussion

In this paper, we construct a reference ontology for maintenance activities based on activity terms used in real-world MWO data. We make use of this reference ontology in an application-level ontology to assess the data quality of MWO records. The use of real-world data impacts the choices made in the construction of this ontology. It is through this lens that the ontology's performance is assessed.

7.1. How did the ontology perform for on the use case?

In Section 6, we demonstrate that the application-level ontology is successful in replicating SME knowledge for every competency question. We use the ontology to compare the activity term appearing in the MWO unstructured text to the logic used by SMEs to assess what work was completed in a work order. In this section, we demonstrate the sucess of the ontology in surfacing data quality issues in the MWOs. We also discuss the assumptions that we made in the construction of this ontology, and possible improvements to limit the impacts of these assumptions in future work.

7.1.1. What worked well?

A summary of the data quality issues identified by the ontology is shown in Table 10. For the 36 records in this pump MWO data set from 2012-2020 there were 20 records (55%) with data quality issues. The two type of data quality issues identified are a) missing data, and 2) inconsistent data. The missing data occurs when there is no maintenance activity term in the MWO unstructured text field. The ontology is able to infer what this missing term might be. As described in Table 11 the ontology is also able to identify inconsistencies between the identified maintenance activity cluster term and the inferred activity classification and route these to an engineer for resolution.

There was 100% agreement between the ontology and the subject matter classifications for maintenance terms. This is because the ontology has been built to replicate the logic shown in Figure 2. The results in Section 6 demonstrates that the ontology and rules can successfully capture the SME knowledge. Note, that both the SME and the ontology used the same assumed values for item and service costs. For other equipment classes (outside of centrifugal pumps) the same assumed values may not be applicable. This is a subject of future work as noted in the following regarding the distinction between repair and replace.

The ability of the ontology to automatically identify data quality issues in these records has the potential to significantly improvement to both measuring data quality through quantifying the number of records with issues, improving data quality by routing affected records to an engineer and/or incorporating these data quality checks when data is entered.

7.1.2. What could be improved?

Distinguishing between repair and replace for corrective actions: For many of the work order classifications, there is insufficient evidence to distinguish between two of the activities (i.e. Repair or Replace and Adjust or Calibrate) without more asset and site specific information. It would be possible to include additional reasoning between repair and replace based on the material cost. If the material cost is equal to the cost of a new unit then the activity would be replace, if not, then it would be a repair. Other options would involve considering the parts assigned to the maintenance work order to see if these parts included, say, a pressure switch (as a replaced component) or not. This would require adding new classes capturing the parts used in the activity. If the parts included a pressure switch then we could reason that a replace activity had occurred. If not, then the activity would be considered a repair. The parts list associated for each MWO number is captured in a separate table in the database. While these improvements could potentially procedure a more generalised approach, this data was not available for the construction of our application-level ontology.

What to do when no activity is identified: In the case of the MWOs with IDs 1, 4, 5, 13, 14, 16, 20, 22, 30, 31, 33 and 35 there is no activity word in the unstructured text field. Instead the problem is described in the text. For example, the MWO with ID 30 says ‘oil leak from housing seal’. Note there are significant labour and maintenance costs incurred ($8632 and $8369). In all cases where there is no activity identified, the work is corrective. An engineer would reason that these activities involve either a repair or replacement. In the application-level ontology, this classification is also made based on the logic in Figure 2. However, since there is no identified activity, we can flag this record for futher examination by an engineer.

Disambiguation of equipment subunits for complex equipment examples (i.e. complex pumps): In this paper we have worked with a use case on a centrifugal pump. In the process industries there are a number of equipment classes that have well defined taxonomies of functional subunits and maintainable items. For example, heat exchangers, compressors, pumps, turbines. CMMS systems in different organisations are set up to use these hierarchies when they have to share data for reference data sets [43]. While the equipment subdivision in Table 5 is fit-for-purpose for standard centrifugal pumps, some specialist pumps might also include a separate pressurized lube oil pump system. Records of maintenance work on these can be difficult for an engineer to disambiguate as there are two pumps in the system and a number of seals, some on the main pump and others on the lube pump. Our focus in this work is on the data quality of the standardised equipment classes which make up the majority of assets in process plants and not on the specialised (and often may be expensive one-of-a-kind) assets. For the latter we recommend that specialist engineers continue to manually review the MWOs.

Disambiguation of Semantic Synonyms for Activities: In this work we have mapped each of the semantic synonyms to one activity class, as shown in Table 6. While this allows us to move forward with supporting data quality work right now we recognise there could be additional information incorporated into this mapping activity and this will be the subject of future collaborative work in both the NLP and ontology communities. For example, we have mapped install to the replace activity cluster. We do not know if it is a like-for-like replacement from the information we have in this case study. We have accounted for future work in this area in the elucidation for the term replace shown in Table 7. This would allow reasoning to be developed for information about the parts that were replaced and the parts installed, when this information is available.

7.2. What insights did we get into maintenance activity classification

One insight gained in this work is the value of determining if maintenance activities are corrective or preventative. This information is contained in the *work order type* field in Table 4. When writing elucidations for the terms in our reference ontology, we found that SME definitions for some activities (derived from descriptions given in international standards), diverge only through the word “periodic". For example, a service activity is a periodic adjust activity. This means that service activities are aligned to a preventative maintenance strategy, and adjust activities are not. This insight informed a key part of the reasoning in the application-level ontology. This insight is consistent with the need for reliability engineers to know if preventative work is being done or not as a check of compliance with maintenance strategy.

Another insight is the importance of an item's state in distinguishing between activity types. To realise this, we link our elucidations to our previous work on Maintenance State [53]. For example, we determined that before a replace activity, an item is in a Degraded State or a Failed State, meaning that the item's "ability to perform its required function" (as per the semi-formal definition) is impaired. After the replace activity, the item is in a Operating State as its required function has been "restored". However, the state of an item is not materially changed throughout an adjust activity. We have used this to distinguish between a repair and adjust activity from an ontological perspective. However, in practice, there is a grey zone between what some would call an adjustment and others a repair. This is worth considering in future work.

In our use case and this application-level ontology, we do not have an item's state information available to us (as condition and performance data is captured in a different data base, not generally in a MWO record). However with the increased emphasis on condition monitoring on assets and process instrumentation of systems this information could be available but would need data from data bases beyond the CMMS. Future work should look at how to use this process and condition data to assist with the decisions being made in this application-level ontology. Digital twins will have to be able to incorporate changes in state resulting from maintenance activities in their models.

7.3. Choice of upper ontology

In Section 4.2 we discussed the limitation we encountered with using BFO for the notion of *functional object* and its associated hierarchies. The concept of an asset having both a functional and a physical hierarchy is central to engineering design and systems engineering. In early design, operation and reliability engineers pay significant attention to the functional perspective, identifying what desired and secondary functions are required and how these functions are fulfilled. Functions are usually associated with desired levels of performance. As the design life cycle proceeds towards manufacturing then physical and structural consideration get increasing attention. Design for maintenance also considers the physical structure of the system, subunits and maintainable items. For ontologies to be applied to industry data, the models and reasoning must be able to accommodate both functional and physical perspectives.

The authors of this paper are part of IOF Maintenance working group and this relationship was the primary reason for selecting BFO as the upper ontology. Our interest is in understanding performance of this upper ontology on a range of use cases requiring reference and application ontologies and the use of real industry data, and using these experiences to encourage further development of BFO. We are also interested in how other ontologies perform and are active also in using the emerging ISO/CD TR 15926-14 ontology for maintenance data [32].

8. Conclusion

There is significant opportunity to use ontologies as an integral part of the data cleaning process for maintenance work order records. This paper focuses specifically on data quality checks on what maintenance activity was performed based on information contained in the MWO. Understanding these maintenance actions on a specific piece of equipment is analogous to a doctor being able to determine all the procedures and prescriptions a patient has had from historical medical records. In doing so, we cannot afford to rely only on the words used by the data generator to describe the *activity*. There are too many different people involved over the years in generating these records and they come from a wide range of backgrounds. As a result, to confirm the activity that was actually performed, engineers have to check if these *activity* words are consistent with data in the other fields of the MWO record. The ontology in this paper performs this check. The results are promising with the reasoning is able to identify several data quality issues in the MWOs. The challenges for the ontology in making a final distinction between classes such as replace or repair and adjust or calibrate for corrective work order types are the same as for the engineers. In both cases additional information is required. The elucidations for each of the seven maintenance activity terms also reveal shortcomings, from an ontology perspective, with the descriptions for some of these *activity* terms in both the ISO 14224 and ISO 15926-4 standards.

The generation of instance data for the reasoning in the application-level ontology is dependent on having a suitable NLP pipeline. Work in NLP for MWOs is an active area of research. While the pipeline used here is state of the art for extracting *activity* and *item* information, we expect additional instance data and relational data to be available soon. This will assist with refining the reasoning of the application-level ontology presented in this paper. The combination of both NLP and ontologies for an industrial use case is an additional contribution to the literature for this paper.

The ontology community has a tendency to regard processes as failures if they perform in a less than perfect way. In the case of improvements in data quality of MWOs the task is so huge and laborious and the consequences of poor data quality so difficult to quantify and yet significant that any automation and standardisation is a worthwhile improvement. This work is a stepping stone to the use of ontologies for cleaning and processing maintenance work orders in industry. A necessary step in the digitisation journey modern industry players are now on.

References

[1] S. Stern, S. Farioli, E. Eisenschmidt, S. Morachioli, C. Kienzler, D. Jeske and N. Schaal, The future of maintenance for distributed assets, *McKinsey White Paper* (2020).

[2] H. Brink, A. Krych, O.R. Cardenas and S. Tiwari, Establishing the right analytics-based maintenance strategy, *McKinsey White Paper* (2021).

[3] M. Bertolini, D. Mezzogori, M. Neroni and F. Zammori, Machine Learning for industrial applications: a comprehensive literature review, *Expert Systems with Applications* **175** (2021), 114820.

[4] B.A. Weiss, M.P. Brundage, Y. Tamm, T. Makila, J. Pellegrino et al., Summary Report on the Industry Forum for Monitoring, Diagnostics, and Prognostics for Manufacturing Operations, *NIST Advanced Manufacturing Series 100-23* (2019).

[5] N. Chilamkurti, T. Torabi and R. Elhdad, Ontology-based framework for maintenance activity analysis and support: A case study for petroleum plant, *International Journal of System Assurance Engineering and Management* **5**(1) (2014), 84–98.

[6] J. Sikorska, M. Hodkiewicz and L. Ma, Prognostic modelling options for remaining useful life estimation by industry, *Mechanical Systems and Signal Processing* **25**(5) (2011), 1803–1836.

[7] M.P. Brundage, T. Sexton, M. Hodkiewicz, A. Dima and S. Lukens, Technical language processing: Unlocking maintenance knowledge, *Manufacturing Letters* **27** (2021), 42–46.

[8] A. Dima, S. Lukens, M. Hodkiewicz, T. Sexton and M.P. Brundage, Adapting natural language processing for technical text, *Applied AI Letters* **2**(3) (2021), e33.

[9] K. Unsworth, E. Adriasola, A. Johnston-Billings, A. Dmitrieva and M. Hodkiewicz, Goal hierarchy: Improving asset data quality by improving motivation, *Reliability Engineering & System Safety* **96**(11) (2011), 1474–1481.

[10] R. Molina, K. Unsworth, M. Hodkiewicz and E. Adriasola, Are managerial pressure, technological control and intrinsic motivation effective in improving data quality?, *Reliability Engineering & System Safety* **119** (2013), 26–34.

[11] G.D. Murphy, Improving the quality of manually acquired data: Applying the theory of planned behaviour to data quality, *Reliability Engineering & System Safety* **94**(12) (2009), 1881–1886.

[12] M. Madhikermi, S. Kubler, J. Robert, A. Buda and K. Främling, Data quality assessment of maintenance reporting procedures, *Expert Systems with Applications* **63** (2016), 145–164.

[13] D. Rajpathak, H. Siva Subramania and P. Bandyopadhyay, Ontology-driven data collection and validation framework for the diagnosis of vehicle health management, *International Journal of Computer Integrated Manufacturing* **25**(9) (2012), 774–789.

[14] M. Hodkiewicz and N. Montgomery, Data fitness for purpose: assessing the quality of industrial data for use in mathematical models, in: *8th International Conference on Modelling in Industrial Maintenance and Reliability, Institute of Mathematics and its Applications, Oxford*, 2014, pp. 125–130.

[15] M. Hodkiewicz and M.T.-W. Ho, Cleaning historical maintenance work order data for reliability analysis, *Journal of Quality in Maintenance Engineering* **22** (2016), 146–463.

[16] T. Sexton, M. Hodkiewicz, M.P. Brundage and T. Smoker, Benchmarking for keyword extraction methodologies in maintenance work orders, in: *Annual Conference of the Prognostics and Health Management Society*, Vol. 10, 2018.

[17] A. Conte, C. Bolland, L. Phan, M. Brundage and T. Sexton, The Impact of Data Quality on Maintenance Work Order Analysis: A Case Study in Historical HVAC Maintenance Work Orders, in: *PHM Society European Conference*, Vol. 6, 2021, pp. 11–11.

[18] T. Bikaun and M. Hodkiewicz, Semi-automated Estimation of Reliability Measures from MaintenanceWork Order Records, in: *PHM Society European Conference*, Vol. 6, 2021, pp. 9–9.

[19] M.V. Ottermo, S. Håbrekke, S. Hauge and L. Bodsberg, Technical Language Processing for Efficient Classification of Failure Events for Safety Critical Equipment, in: *PHM Society European Conference*, Vol. 6, 2021, pp. 9–9.

[20] D. Rajpathak and R. Chougule, A generic ontology development framework for data integration and decision support in a distributed environment, *International Journal of Computer Integrated Manufacturing* **24**(2) (2011), 154–170.

[21] M. Navinchandran, M.E. Sharp, M.P. Brundage and T.B. Sexton, Discovering critical KPI factors from natural language in maintenance work orders, *Journal of Intelligent Manufacturing* (2021), 1–19.

[22] IEC, Dependability management – Maintenance and maintenance support, Standard, AS IEC 60300.3.14, International Electrotechnical Commission, Geneva, Switzerland, 2016.

[23] IEC, IEC 60812 Failure modes and effects analysis (FMEA and FMECA), Technical Report, IEC, Geneva, Switzerland, 2018.
[24] SAE International, A Guide to the Reliability-Centered Maintenance (RCM) Standard, Standard, SAE JA1012, SAE International, 2018.
[25] M. Hodkiewicz, S. Lukens, M.P. Brundage and T. Sexton, Rethinking maintenance terminology for an Industry 4.0 future, *International Journal of Prognostics and Health Management* **12**(1) (2021).
[26] ISO/IEC, ISO/IEC 21838-2:2021 Information technology — Top-level ontologies (TLO) - Part 1: Requirements, Technical Report, ISO/IEC, Geneva, Switzerland, 2021.
[27] ISO/IEC, ISO/IEC 21838-2:2021 Information technology — Top-level ontologies (TLO) — Part 2: Basic Formal Ontology (BFO), Technical Report, ISO/IEC, Geneva, Switzerland, 2021.
[28] B. Smith, F. Ameri, H. Cheong, D. Kiritsis, D. Sormaz, C. Will and J.N. Otte, A First-Order Logic Formalization of the Industrial Ontology Foundry Signature Using Basic Formal Ontology, in: *10th International Workshop of Formal Ontology Meets Industry (FOMI)*, Graz, Austria, 2019.
[29] A. Matsokis, H.M. Karray, B. Chebel-Morello and D. Kiritsis, An ontology-based model for providing semantic maintenance, *IFAC Proceedings Volumes* **43**(3) (2010), 12–17.
[30] M.H. Karray, F. Ameri, M. Hodkiewicz and T. Louge, ROMAIN: Towards a BFO compliant reference ontology for industrial maintenance, *Applied Ontology* **14**(2) (2019), 155–177.
[31] M. Hodkiewicz, C. Woods, E. Low and F. Ameri, Towards a Reference Ontology for Maintenance Work Management, in: *Workshop Proceedings of the Conference on Enterprise Interoperability and Applications – Interoperability in the era of Artificial Intelligence (I-ESA 2020)*, 2020.
[32] M. Hodkiewicz, J.W. Klüwer, C. Woods, T. Smoker and T. French, An ontology for reasoning over engineering textual data stored in FMEA spreadsheet tables, *Computers in Industry* **131** (2021).
[33] D. Rajpathak, Y. Xu and I. Gibbs, An integrated framework for automatic ontology learning from unstructured repair text data for effective fault detection and isolation in automotive domain, *Computers in Industry* **123** (2020), 103338.
[34] D.G. Rajpathak, An ontology based text mining system for knowledge discovery from the diagnosis data in the automotive domain, *Computers in Industry* **64**(5) (2013), 565–580.
[35] J.I. Single, J. Schmidt and J. Denecke, Knowledge acquisition from chemical accident databases using an ontology-based method and natural language processing, *Safety Science* **129** (2020), 104747.
[36] D.P. Lupp, M. Hodkiewicz and M.G. Skjæveland, Template libraries for industrial asset maintenance: A methodology for scalable and maintainable ontologies, in: *CEUR Workshop Proceedings*, Vol. 2757, Technical University of Aachen, 2020, pp. 49–64.
[37] L. Nagy, T. Ruppert and J. Abonyi, Ontology-based analysis of manufacturing processes: lessons learned from the case study of wire harness production, *Complexity* **2021** (2021).
[38] Y. Gao, C. Woods, W. Liu, T. French and M. Hodkiewicz, Pipeline for machine reading of unstructured maintenance work order records, in: *Proceedings of the 30th European Safety and Reliability Conference and 15th Probabilistic Safety Assessment and Management Conference*, 2020.
[39] T.B. Sexton, M.P. Brundage et al., Nestor: A tool for natural language annotation of short texts, *Journal of Research of the National Institute of Standards and Technology* **124** (2019).
[40] M. Stewart, W. Liu and R. Cardell-Oliver, Redcoat: a collaborative annotation tool for hierarchical entity typing, in: *Proceedings of the 2019 Conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP): System Demonstrations*, 2019, pp. 193–198.
[41] T. Sexton, M. Hodkiewicz, M.P. Brundage et al., Categorization errors for data entry in maintenance work-orders, in: *Annual Conference of the Prognostics and Health Management Society*, Vol. 11, 2019.
[42] A.T. James, O. Gandhi and S. Deshmukh, Knowledge management of automobile system failures through development of failure knowledge ontology from maintenance experience, *Journal of Advances in Management Research* (2017).
[43] ISO, ISO 14224: Petroleum, petrochemical and natural gas industries – Collection and exchange of reliability and maintenance data for equipment, Standard, ISO14224:2016, International Organization for Standardization, Geneva, Switzerland, 2016.
[44] ISO, ISO 15926-4:2019 Industrial automation systems and integration—Integration of life-cycle data for process plants including oil and gas production facilities - Part 4: Initial reference data, Technical Report, ISO, Geneva, Switzerland, 2019.
[45] SINTEF, Offshore Reliability Data Handbook, Technical Report, Høvik: Det Norske Veritas, 2015.
[46] M. Rausand and K. Øien, The basic concepts of failure analysis, *Reliability Engineering & System Safety* **53**(1) (1996), 73–83.
[47] M. Stewart, W. Liu, R. Cardell-Oliver and R. Wang, Short-Text Lexical Normalisation on Industrial Log Data, in: *2018 IEEE International Conference on Big Knowledge (ICBK)*, 2018, pp. 113–122. doi:10.1109/ICBK.2018.00023.
[48] M.H. Karray, B. Chebel-Morello and N. Zerhouni, A formal ontology for industrial maintenance, *Applied Ontology* **7**(3) (2012), 269–310.
[49] UWA NLP-TLP Group, The University of Western Australia: Natural & Technical Language Processing Group, 2021. https://nlp-tlp.org/.
[50] T. Bikaun, T. French, M. Hodkiewicz, M. Stewart and W. Liu, LexiClean: An annotation tool for rapid multi-task lexical normalisation, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2021, pp. 212–219.
[51] SAE, Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes, Standard, JA1011-2009, SAE International, 2009.
[52] International Organization for Standardization, ISO 15926-14: Industrial automation systems and integration – Integration of life-cycle data for process plants including oil and gas production facilities – Part 14: Data model adopted for OWL 2 Direct Semantics, Community Draft, 2019.
[53] C. Woods, M. Selway, M. Hodkiewicz, F. Ameri, M. Stumptner and W. Sobel, On the Notion of Maintenance State for Industrial Assets, in: *CEUR Workshop Proceedings*, Vol. 2969, Rheinisch-Westfaelische Technische Hochschule Aachen* Lehrstuhl Informatik V, 2021.
[54] H. Deutsch and P. Garbacz, Relative identity (2002).

[55] C. Will, Industrial Ontology Foundry (IOF) Core, 2021. https://github.com/NCOR-US/IOF-BFO/tree/IOF-Core-2020.

Supplementary materials

Table 11

Maintenance activity verbs extracted from 6000 collaboratively annotated maintenance work orders. **Bold** indicates a maintenance activity and no font indicates a supporting activity.

access	clear	drive	instrument	overload	**recharge**	reschedule	shutdown	**unblock**
activate	clip	drop	invert	pack up	reclaim	**reseal**	sign	**unbog**
add	close	duplicate	invest	**patch**	**reclamp**	**reseat**	spray	unload
adjust	coil	energize	**investigate**	perform	**reconnect**	**re-secure**	stage	upgrade
align	collect	erect	isolate	pop out	**rectify**	**reset**	start	**verify**
amend	commission	extend	label	**position**	redesign	**re-splice**	start-up	weld
analyse	communicate	fabricate	laminate	prepare	**refill**	**re-tension**	**straighten**	
apply	compile	face	launch	prime	**refit**	**retest**	supply	
approve	complete	fasten	launder	print	**refurbish**	**retorque**	support	
assemble	conduct	**fault-find**	lift	process	reinforce	**retrack**	survey	
assess	confirm	**fill**	load	program	**reinstall**	return	suspend	
assist	connect	**find**	locate	project	**reinstate**	review	**swap**	
attach	**correct**	fit	look	pull	reject	**reweld**	take up	
attend	create	**fix**	**lubricate**	**purge**	release	**rework**	**test**	
audit	cut	follow-up	**maintain**	push	**reline**	rip	**thermography**	
buff	deflect	**free-up**	make	put	relocate	**rotate**	**thickness test**	
build	deisolate	generate	manufacture	quote	**remount**	run	**tighten**	
bypass	demobilise	**grease**	**measure**	raise	**remove**	**sample**	**tilt**	
calibrate	design	**guide**	mobilize	**re-engage**	reopen	**scan**	**top up**	
call	**detect**	heat	modify	**reactivate**	**repair**	scope	**track**	
carry	develop	hold	**monitor**	read	**replace**	scrap	train	
centralise	**diagnose**	**hose out**	move	**readjust**	report	secure	tramp	
change	disable	**identify**	obtain	**realign**	**re-position**	sequence	transfer	
change out	**discharge**	improve	open	**reattach**	request	**service**	transition	
change over	disconnect	insert	operate	**rebuild**	reroute	set	trial	
charge	**drain**	**inspect**	order	**recalibrate**	rerun	set-up	trim	
check		install	overhaul				**troubleshoot**	
clean		instigate					**tune**	