Daniel Sonntag, Sonja Zillner, Patrick van der Smagt, András Lőrincz

Overview of the CPS for Smart Factories Project: Deep Learning, Knowledge Acquisition, Anomaly Detection & Intelligent User Interfaces

Version 7

Keywords: cyber-physical systems, smart factories, deep learning, knowledge acquisition, anomaly detection, intelligent user interfaces

DFKI, Daniel Sonntag, sonntag@dfki.de
Siemens, Sonja Zillner, zillner@siemens.com
fortiss & TUM, Patrick van der Smagt, smagt@brml.org
ELTE, András Lőrincz, lorincz@inf.elte.hu

Book Title: Industrial Internet of Things: Cyber-Manufacturing Systems, CPS-Book (Eds. Jeschke, Brecher, Song, Rawat)

Contents

1 Introduction **3**

2 Technical Infrastructure **4**
2.1 Deep Learning 4
2.2 Knowledge Acquisition 5
2.3 Anomaly Detection 6
2.4 Intelligent User Interfaces 9

3 Use Cases **12**
3.1 Industrial Robots and Anomaly Modeling 12
3.2 Anomaly Treatment in the Steel Domain 13
3.3 Outlook: Anomaly Detection in the Energy Domain 16

4 Acknowledgments **18**

Abstract

Industry 4.0 factories become more and more complex with increased maintenance costs. Reducing costs by cyber-physical (CP) controllers should ensure the commercialization of the *CPS for smart factory* project results. We implement multi-adaptive CP controllers in the following domains: industrial robot arms, car manufacturing, steel industry, and assembly lines in general. The main objective is to implement such controllers for these application domains and let the industry partners provide feedback about the cost reduction potential. In this paper, we describe the technical infrastructure including deep learning and knowledge acquisition submodules, followed by anomaly detection modules and intelligent user interfaces in the IoT (Internet of Things) paradigm. In addition, we report on three concrete use case implementations of industrial robots and anomaly modeling, knowledge management and anomaly treatment in the steel domain, and anomaly detection in the energy domain.

1 Introduction

A tight cooperation of automation and IT vendors should enable sustainable business models supporting the European manufacturing sector to manage its increasingly complex, inter-organizational production networks and align them efficiently with global supply chains. The individual components will be ready as products or as input for product development. Innovation is supported by evaluated business models and concrete examples for customer business cases.

This should be realized by integrating a platform which uniquely combines cross-enterprise event management (anomaly treatment via deep learning, knowledge management via a semantic portal, intelligent user interfaces) with digital product memory technology and smart object virtualization.

In this chapter, we report on our three milestones of the *CPS for smart factories* activity at EIT Digital, funded by the EU.[1] The three milestones can be summarized as

1. CPS Knowledge engineering: understanding the formal requirements of application cases and their formalization;
2. Implementation of software modules and tuning formal models and rules to test scenarios of anomaly detection in physical environments. This includes functional programming with deep learning capacity, and ontology creation/manipulation/extension via a semantic portal infrastructure and intelligent user interfaces [28];
3. Transfer of modules into industrial settings, first evaluations, and business modeling.

[1] See http://dfki.de/smartfactories.

2 Technical Infrastructure

The over-arching challenge to address is to combine cyber-physical system (CPS) safety and performance. CPS is to be understood as a network of interacting elements with physical input and output, forming a system of collaborating computational elements controlling physical entities such as Industry 4.0 factories. While addressing safety challenges, the outcomes include models of the behavior of loops with human operators, in particular how to ensure safety. Beyond failures in the robotic system, humans can also make mistakes, and thus a special desired outcome is a model which accounts for humans as producing anomalies by reacting to predictable maintenance tasks and unpredictable events. Technical advancements include, most notably:

1. a GPU-based deep learning machine learning infrastructure for anomaly treatment and data mining;
2. a smart factories knowledge portal infrastructure for an anomaly instance base (knowledge acquisition and management);
3. model-based predictions with anomaly detection followed by workflow management, including real-time verification and, possibly, machine learning fostering earlier anomaly detection;
4. intelligent user interfaces for expert knowledge acquisition, human behaviour input, and human-robot interaction by using, e.g., vision sensors.

These components are built into an architecture and are to be extended with the characterization of a human collaborator who is also *in the loop* and may also exhibit anomalous behavior. Cyber-physical systems are implemented in human environments. The software/hardware outcome package consists of an anomaly management system, including controllers for smart factories. We focus on both open and closed-loop controllers in the robot domain and reporting/maintenance domain in manufacturing.

The work plan for future Industry 4.0 factories comprises business models for Industry 4.0 technology in order to tackle “unmodeled” anomalies that need to be counteracted. That may happen in diverse ways: the first task is the detection and clustering of the anomaly, followed by modeling by means of human expert domain knowledge, and finally, the computer-assisted optimization, including the extension of the ontology or anomaly dictionary and the related (automated) cost-saving workflow management.

2.1 Deep Learning

Based on recent research results, we exploit deep neural networks for the representation and dimensionality reduction of complex sensory data. Neural networks are often understood as the universal approximators studied in the late 1980s and early 1990s, having a few layers of hidden sigmoidal units. It has been shown that the training of such shallow networks is NP-complete [17]. Deep networks (or deep belief networks) have the structure of the above-mentioned neural networks but a larger number of hidden layers; typically, this may go beyond 100 [10]. Efficient methods have been introduced recently, see, e.g., [14] and [25]. The new generation of deep neural networks can do more than associating outputs to inputs; they can work in reverse and can generate inputs from representations, see [5] for a survey. Such networks are called autoencoders. They have representational and generalization capabilities far beyond that of many others. Generalization capabilities are excellent, but still: the larger and richer the database, the better is the performance for most application cases. The resurrection of neural networks was caused by three important factors:

- the development of deep learning, including the solutions for the vanishing gradient problem by Restricted Boltzmann Machines [14, 5, 25];
- the increase of data set sizes via crowdsourcing;
- the use of graphics processors (GPUs) in computation, leading to processing speed increase of up to two orders of magnitude. This increase allows for serious hyper-parameter searches even in large data sets, eventually leading to better optimization.

It must be stressed that there is still an important dichotomy between neural network and Bayesian machine learning. For a large part, Bayesian analysis does not apply to nonlinear neural networks and a rigorous mathematical analysis of methods or results is not yet within our possibilities. This used to be the case for deep learning, too, but recent developments provide provable bounds for some networks types [2]. Probabilistic neural networks have been developed [4, 3] and are in wide use. Furthermore, variational approximations that exploit the autoencoder concept gained momentum recently [18, 16].

2.2 Knowledge Acquisition

Production controllers and their contextualization demand for a *real-time* semantic layer for, e.g., the assembly lines in automotive factories, the steel production domain, or the energy domain with smart meter and smart grid analytics applications.

In order to comply with the underlying logic of the daily business operations, we rely on a dedicated semantic model supporting longitudinal access across heterogeneous data sources. The semantic model, including the semantic portal implementations and human-in-the-loop knowledge acquisition, will establish the basis for seamless data integration of all production applications.

For the steel production use case, a semantic mediawiki[2] architecture has been implemented. A major goal of this architecture was to combine static facility models to be stored in an RDF triple store and made accessible to the user via semantically enriched MediaWiki pages with dynamically executed business process models in a largely seamless way. In this specific system architecture of

[2] https://semantic-mediawiki.org

the hot rolling mill (steel usecase) there are two sources of information: a digital pen (or pen based interaction on a corresponding smart phone application) and the Object Memory Server (OMS) [13] for accessing an OMS memory that can be stored in simple and cheap RFID labels. This information is then shared with the semantic mediawiki to provide sensor data from the hot rolling mill parts or the production line parts in real-time. The goal of this knowledge portal development is the realization of a situation-specific adaption of production steps, i.e., process parameters, because of certain potential anomaly events during the manufacturing process. The dynamic behavior of smart machines was realized by a knowledge-based decision making component. The component decides on the actions to be taken in the production process based on information about product variants and machine capabilities described in specifically designed ontologies.

As a foundation of semantic product memories for the internet of things and this smart factories project, see [35]. The development of low-cost, compact digital storage, sensor and radio modules allows us to embed digital memories into products to record those anomaly key events. Such computationally enhanced products can perceive and control their environment, analyze their observations, and communicate with other smart objects and human users. The RFID and semantic portal infrastructure supports the interaction with digital product memories [19] and controlled interaction with digital product memories [12] during an assembly process. In addition, user input data is forwarded to the corresponding servers on which additional handwriting and gesture recognition (active user input) or sensor checking procedures (automatic passive sensor input) are executed. These results (e.g., documented anomalies) are stored in the respective XML documents of the semantic MediaWiki (see Figure 1).

2.3 Anomaly Detection

Models are required for describing and controlling the dynamic behavior of nonlinear plants. Typically, such models are not sufficiently rich, especially if the plant has many degrees of freedom or high-dimensional sensors, or when it is embedded into a complex environment like robotic systems, intelligent vehicles, or any other modern actor-sensor system that we depend on. In such cases, the quality of fault detection deteriorates: too many false positives (i.e., false alarms) make the fault detection useless, while too many false negatives (i.e., unobserved faults) may harm the system. Rather than fully trusting incomplete models, we have put forth a methodology which creates a *probabilistic vector time series model* of the system from the recorded data and detects outliers, also called anomalies with respect to this learned model. This type of detection is notoriously difficult as it is an ill-posed problem. First, the notion of anomaly strongly depends on the domain. Then, the boundary between “normal” and “anomalous” might not be precise and might evolve over time. Also, anomalies might appear normal or be obscured by noise. Finally, collecting anomalous data is very difficult, and labeling them is even more so [9].

Two observations are important to make: (i) anomalies are sparse by their

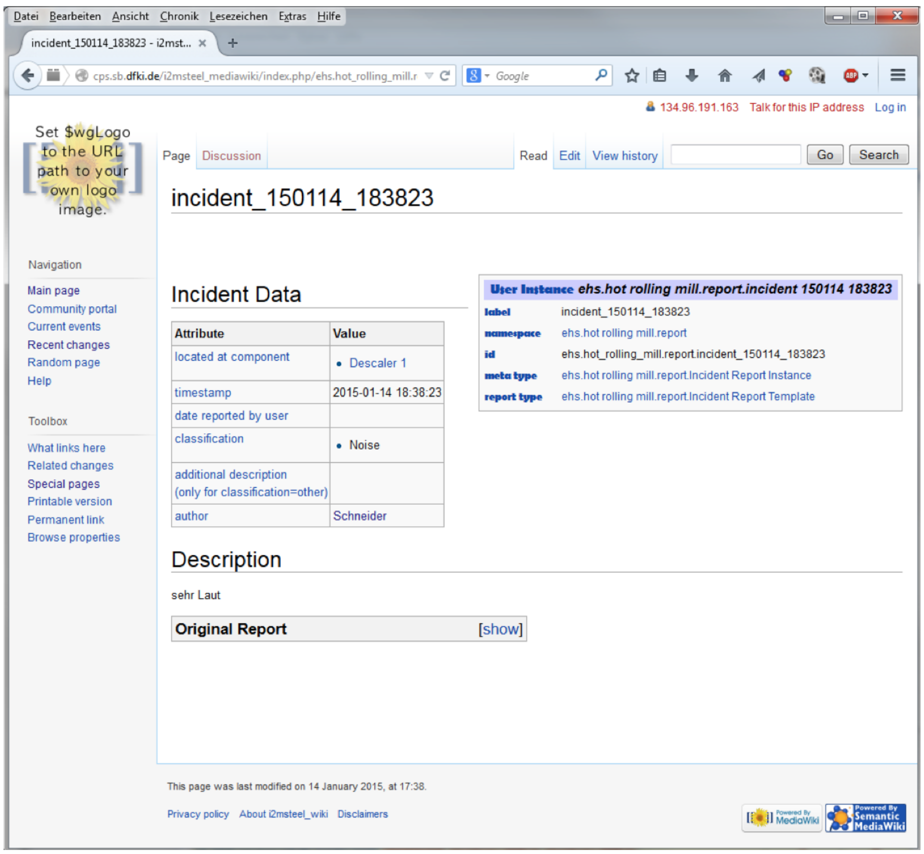

Figure 1: Example of an automatically generated incident report (anomaly)

Figure 2: Frame series of the assistant hitting the robot arm

very nature, and (ii) in a high-dimensional real-world scenario, it will not be possible to rigorously define “normal” and “anomalous” regions of the data space. We therefore designed an unsupervised approach together with a data collection machinery by using either human interactions (Figure 2) or machine generated “anomalies” (Figure 3.)

Figure 2: In one step of the data generation process, a probabilistic vector time series model of the system’s data is created, and (patterns of) samples that do not fit in the model were conjectured to be the anomalies. We found that searches for anomalies can be made robust by means of Robust Principal Component Analysis [8] if it is combined with group fused Lasso techniques [7] and sparse event filtering [22].

Figure 3: The seven measured joint angle trajectories are plotted in different colors. The arrival of a new desired configuration for the robot arm is visualized by a thin vertical black line. A trajectory segment from the current to the desired configuration is computed and executed. Upon arrival at the desired configuration, a new desired configuration is sampled. In each of the 10 segments there is a probability of 15% that an anomaly is introduced, visualized by a vertical, light colored bar matching the color of the joint the anomaly was induced on. In this example, the anomaly in joint w1 is clearly visible as an indent in the sequence ranging from about 14 s to 15 s. For an anomaly-free sample, the plot would look the same except that there would be no “command” in the sequence.

Although sparse methods are efficient, they are also slow. In order to overcome this critical issue, we developed an architecture that fits the deep learning

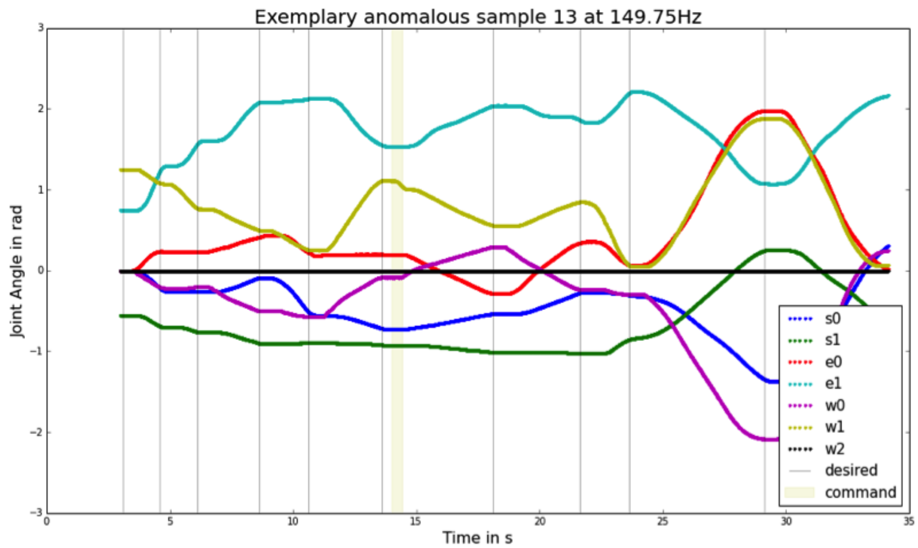

Figure 3: Exemplary anomalous sample.

scheme that we call “Columnar Machine” [20]; this machine can take advantage of the group structure (if present) as in the above-mentioned robust anomaly detection scheme.

Both workflow management and verification require a modular structure and goal-oriented optimization techniques. We developed a theoretical framework for cyber-physical systems with learnable stochastic models of the environment for risk management [32]. The framework meets the constraints of functional programming [21], a desired feature in software development, testing, and verification.

2.4 Intelligent User Interfaces

Internet of Things (IoT) is mainly about connected devices embedded in our everyday environment. Typically, ”interaction” in the context of IoT means interfaces which allow people to either monitor or configure IoT devices. Some examples include mobile applications and embedded touchscreens for control of various functions (e.g., lights or control buttons) in environments such as smart factories. In our application cases, humans are an explicit part of the scenario. Traditional graphical interfaces often lead to a clumsy co-existence of human and IoT devices (consider a tablet for remote-controlling a robot arm). Thus, there is a need to investigate what kinds of interaction techniques could provide IoT to be more human-oriented, what role automation and interaction has to play, and how human-originated data (sensor data for physiological computing) can be used in IoT [33]. Figure 4 shows our intelligent user interface IoT architecture including the semantic mediawiki, the industrial Baxter robot, and the object-

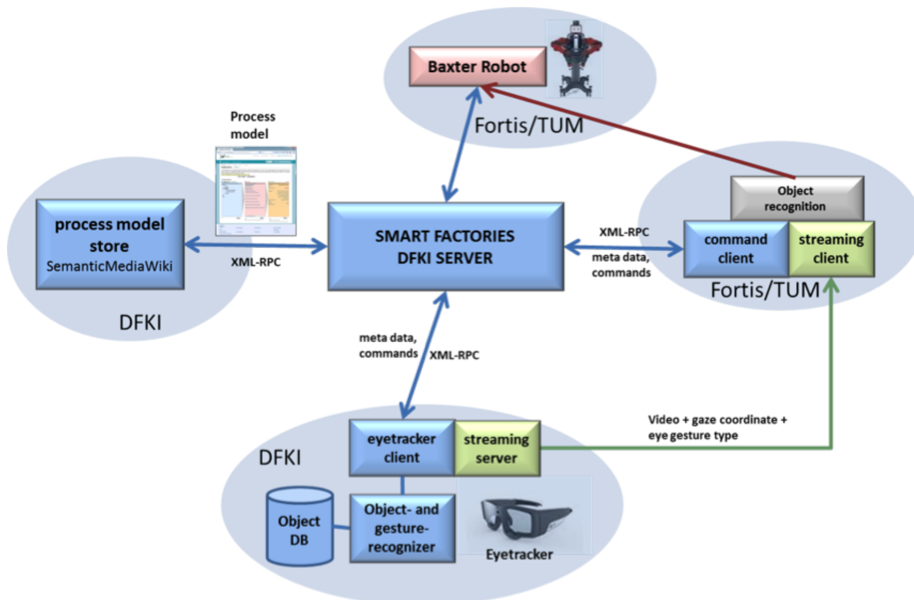

Figure 4: Intelligent User Interface Architecture

detection based anomaly detection.

In many scenarios in modern work practices where robots improve traditional industrial work flows, the co-operation of robots with human workers plays a central role. In this project we focus on direct human-robot-interaction in scenarios where humans and robots interact with the same object (workpiece or tool). Thereby, we address the mutual identification of an object through the human and the robot. It is easy to inform the user about which object the robot is addressing (e.g., by pointing, synthesized speech, explicit action). One of the challenges is to inform the robot about the object the human is attending to without interrupting and affecting the execution of the manual task of the human worker. Figure 5 illustrates the scenario. For the purpose of detecting human intentions automatically, and support activity recognition in general, we developed an eye-tracking system capable of analyzing human eye movements and interpreting specific movement patterns and fixations as eye gestures that were sent to the robot's control system.

To build a base system for experiments leading to such a functionality, it was necessary to develop an operating software for a commercial eye-tracking hardware. Available standard software lacks important features like network transport of gaze points and objects the user gazes at. Also the support of cloud based databases for visual object features that are important in co-operative industrial applications was typically missing in commercial eye-tracking systems. Therefore the distributed service architecture of the project was extended by components around an object- and eye-gesture-recognition service. This system

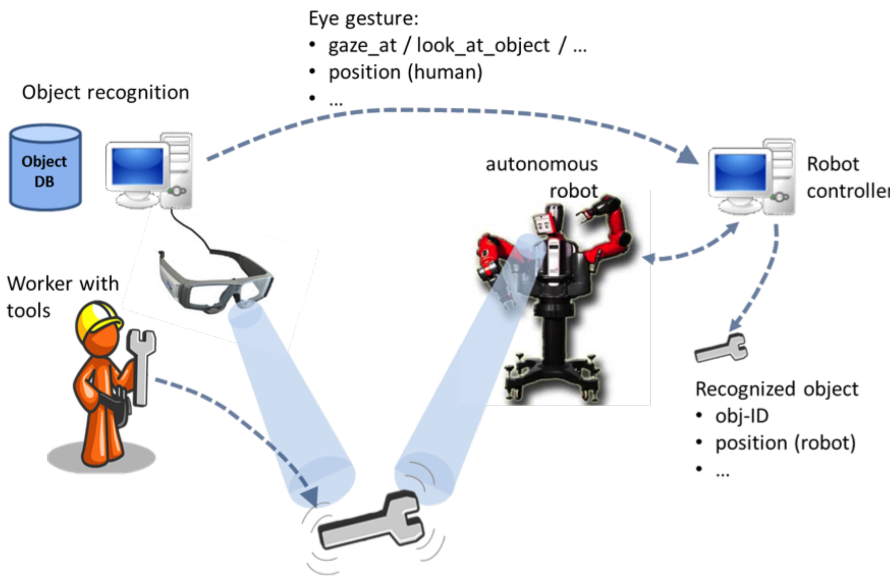

Figure 5: Overview of the co-operation of robots with human workers

uses an object database for storing visual feature vectors of given objects, analyzes the user's eye movements, extracts the gaze coordinates, and identifies different eye gestures. The extracted information is then sent to a client via an efficient low level network protocol. The lower part of figure 4 visualizes the described components of the system architecture based on [29, 30, 27, 31].

In order to allow for vision-based object and activity recognition, we implemented a 3D video annotation tool to provide supervised learning material for the deep learning and anomaly detection infrastructure. [11] present an overview of the state-of-the-art in image and video annotation tools. Two new directions are prominent for Industry 4.0 CPS: first, recent work leverages on highly capable devices such as smartphones and tablets that embrace novel interaction paradigms, for example, touch, gesture-based or physical content interaction [26]; we generalize this to multimodal multisensor annotation in the smart factory context. Second, our own previous development LabelMovie, a semi-supervised machine annotation tool with quality assurance and crowdsourcing options, has been opted for videos (spatio-temporal annotation) [23]. The annotation tool provides a special graphical user interface for multimodal multisensor data and connectors to commercially available sensor systems (e.g., Intel RealSense F200 3D camera, Leap Motion, and Myo). The collected videos are used to learn hand movement for activity recognition in the joint interaction space of a human and a robot.

3 Use Cases

The tight cooperation of automation, cyber-physical systems in production engineering, and IT vendors should enable sustainable business models supporting the European manufacturing and smart grid sector to manage its increasingly complex, inter-organizational production networks and align them efficiently with global supply chains. The individual components of the technical architecture should be made ready as service products or as input for product development. Innovation is supported by evaluated business models and concrete customer business cases, e.g., online anomaly-detection methods in a welding-based production scenario. Recent technological developments include breakthroughs in object memories (smart meter data), big data analysis, and controller software.[3] [4]

We include new methods for finding and treating anomalies such as deep neural networks and the semantic portal that can host instance bases of anomalies.

3.1 Industrial Robots and Anomaly Modeling

We exploit online anomaly-detection methods in production scenarios. Quality control in a production scenario is an important attainable goal, and our methodologies carry the potential of obtaining that by combining deep learning with sparse representations [22].

Human-robot interaction requires the detection, interpretation, and prediction of human body movements within context, including hand, arm, head, face, and eye movements that reveal information about the manipulations and thus about the ongoing activity, including the intentions. We have been developing a novel database for hand pose tracking and exploit deep-learning methods for the estimations. The tool is under testing and the data set can be extended if needed. In this project, we use a hand model (libhand, [34]), 3D cameras of different kinds, and a SmartGlove[5]. After initial evaluation it can be said that only the combination of SmartGlove and libhand suits our goals (precise measurements). Data collected with SmartGlove are transferred to libhand and different 3D views are collected. In current work, we explore robotic motion together with facial expression and human body distance (which is required for safety reasons). A facial expression estimation tool is also available to us, based on [15]. Head pose estimation is very precise, the gaze estimation tool has only about 2 degrees of uncertainty, but requires a high-resolution input video.

In many scenarios in modern work practices where robots improve traditional industrial work flows, the co-operation of robots with human workers plays a central role. For human-robot interaction, we developed a Unity3D serious game in order to model realistic scenarios. This game examines human behavior in tasks of divided attention. The game was implemented for the 3D virtual reality

[3] http://sites.tcs.com/big-data-study/return-on-investment-in-big-data/
[4] http://www.greentechmedia.com/research/report/the-soft-grid-2013
[5] http://www.neofect.com/en/smartglove

headset Oculus Rift, equipped with hand pose estimation 3D camera and the SMI tool for gaze direction estimation. Serious games scripts can be run. The game of the use case was carefully chosen out of many dozens of other games, the main point being to develop a user model, where WCET (worst case execution time) distribution can be estimated. In this game, the need for workflow management becomes straightforward, since attended regions and possibly intentions can be estimated from gaze, together with the registration of the simple manipulation tasks. This is a model for a controller room, where divided attention can be measured under various conditions, including adjustable stress levels. In turn, these games provide a high quality model for CPS with human-in-the-loop problems. A further goal beyond the quantification of the behavior of the human participant is to find methods for proper robotic help in case of anomalies. A number of experiments have been conducted with 10 people and in about 10 sessions. Estimation of WCET is in progress. As a result, this game goes beyond the problem of human-robot collaboration.

3.2 Anomaly Treatment in the Steel Domain

The main focus of this use case scenario is to enable a seamless integration of production and maintenance processes in the context of anomaly treatment. Proper maintenance of industrial plants is of high relevance. It helps to significantly reduce operating costs as well as to improve productivity of the plant operations and the quality of the product. The overall objective of a plant maintenance management system is to ensure the reliability of a plant (component) to perform its functions. Thus, maintenance is seen as any activity that is carried out on a plant or respectively component of a plant in order to ensure that this plant or component of a plant continues to perform its intended function.

However, as of today, the integration of production and maintenance processes and know-how is only addressed and realized in a very limited way. In the past, when the machines have not been that automated, complex and connected, employees from the production site included maintenance task into their daily routines. Only in situations when the handling of the identified failure exceeded their own expertise, external maintenance supports were requested.

Today, with the increase in automatization and digitalization of plants, more and more monitoring and maintenance applications are available. Those monitoring applications are primarily designed to track single and isolated components or parts. Other frequently used maintenance routines are predictive maintenance application servicing single parts or components of plants, such as condition-based monitoring applications for individual plant operations or plant components. In general, those techniques are complemented with preventive maintenance strategies. For example, on a predetermined periodical basis, components of the plant are taken off-line in order to inspect them. Based on the inspection result, repairs are made and the components are put back into operations or the affected components are replaced. Thus, due to the sheer complexity of the underlying processes and operations, this leads to the situation that the particular employees having the highest experience with handling the machines and plant components are no longer actively involved in the maintenance process. In sum, although various maintenance applications and efforts are accomplished on different levels, several shortcomings can be observed:

1. many different monitoring applications, such as predictive maintenance applications, provide important insights about plant components, but do not produce insights covering the comprehensive perspective on the plant performance. The maintenance is focused on local aspects but ignores the plant performance as a whole;

2. the knowledge and expertise of production employees most experienced in handling the plant components are no longer actively involved in the maintenance process;

3. the semantic knowledge about the structure and the basic principles of the plant is not incorporated into the maintenance processes. In particular, the semantics about how plant components are operating and how they are connected with each other as input for the interpretation of local maintenance observations in a global scale is neither available nor used.

In order to overcome the described shortcomings, the general idea of the extended business use case is to seamlessly align human-generated expert know-how with machine-generated maintenance know-how in a semantically consistent manner in order to significantly improve the analytics-based maintenance applications (Figure 6). The main contribution of the described technical infrastructure for selected extended business scenario is

1. to seamlessly align the local perspective of the large number of monitoring applications, such as predictive maintenance applications, as well as the historical data of the various preventive maintenance strategies into one global and coherent perspective;

2. to establish means for the seamless integration and processing of expert knowledge in the production field;

3. to incorporate the structural knowledge about the plant and its operations by means of a semantic model covering various levels such that the general representation of all necessary (pre-processed) data sources becomes possible; and

4. to use this integrated data source as input for analytics applications aiming to produce new valuable insights as well as to trigger automatically recommended actions for improved overall plant maintenance processes.

As already indicated above, existing plant monitoring and maintenance application can be classified as follows [6]:

– *Breakdown maintenance applications*; i.e., applications that help to fix (the component of) a plant when it is broken. This maintenance routine is reactive and is only executed when plant equipment needs to be repaired. It is neither based on an underlying routine maintenance task, nor on a scheduled maintenance strategy.

- *Preventive maintenance applications*; i.e., applications that are based on fixed maintenance schedules in order to replace the affected components of a plant.
- *Predictive maintenance application*; i.e., applications for the condition-based monitoring of the plant operations / components.
- *Proactive maintenance applications*; i.e., application that concentrate on the proactive monitoring of plant operations /components enhanced by the correction of root causes to component failures.

The maintenance application rely on a high degree of automization and digitalization. To the best of our knowledge, there exists no application that allows to seamlessly integrate and process human-based and factory-based monitoring data to improve the overall performance. There are related working papers highlighting the need to align maintenance and production processes. However, besides highlighting the problem scope, those works, e.g., [1], do not provide any concrete solution on how to address these shortcomings. In addition, there exist several research approaches that investigate the use of semantics and ontologies for improving maintenance processes and intelligent fault diagnoses. For instance, [36] introduces an ontology-based reasoning framework for intelligent fault diagnosis of wind turbines. However, this approach neither covers the seamless integration of human-generated data, nor the seamless integration into related business processes. [24] introduces an approach for modeling the semantics of a failure context in order to improve the maintenance support for mobile actors. Although this approach makes use of semantics to formally describe the failure context, the overall applications scenario does not focus on the seamless alignment of human- and machine-generated know-how.

The overall idea of our use case scenario is to establish an application that allows the seamless integration, alignment, processing and analyzing of machine and human-generated monitoring data in order to produce new insights with business value. The machine-generated data originates from the various condition-based monitoring applications or the data repositories of the plan. The human-generated data is captured by semantic-based / intelligent user interaction applications. In order to realize the extended business scenario, several components are required (Figure 6):

The *plant data repository* (our Semantic MediaWiki) encompasses all data sources produced and stored in the context of the production process. For instance, the plant data repository collects any historical information about the accomplished production processes (i.e., all accomplished transformations). The data sources can be distributed along the process production chain. A set of *monitoring application(s)* is needed; these continuously measure the condition of component in order to assess whether it will fail during some future period. The

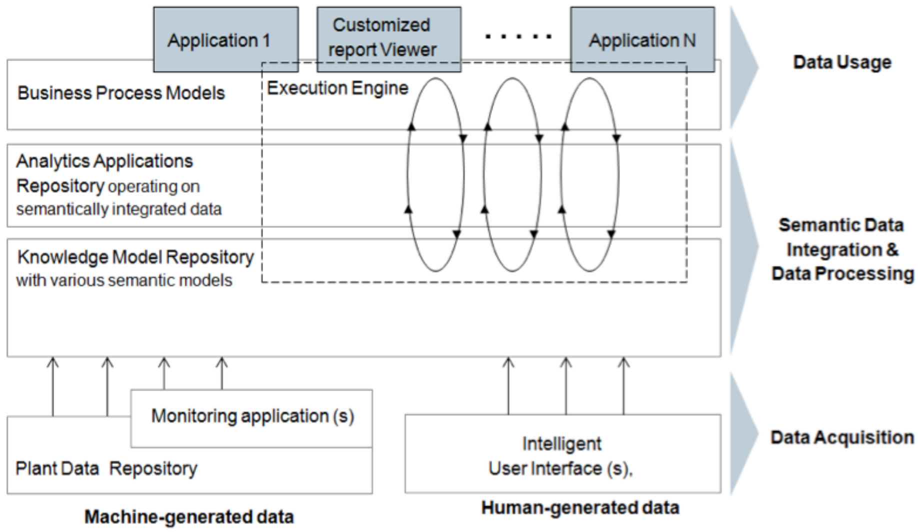

Figure 6: Overview of the steel domain business scenario

data collected in general focus on tracking the condition of particular features, such as vibration analysis, infra-red thermographs, ultrasonic detection, etc. Although the recorded data can be utilized to determine the condition of an isolated component in order to decide about any necessary repairs, the recorded data about possible anomalies will be seamlessly aligned within the adapted incident-report model to establish the basis for aligning human and machine-produced monitoring data for improved plant performance. An *intelligent user interfaces* establishes means to acquire human-produced data in an efficient manner. This semantic data acquisition component, which can be realized as a smart pen application, is aligned with the underlying working routine of the experts in order ensure ease-of use. In addition, the intelligent user interface ensures that all data is captured in semantically annotated form. In general, this is realized by determining the underlying context of the user-input and expressing it with corresponding semantic terminology. In this implementation of the use case scenario, we are using a smart pen application in combination with a customized incident / anomaly report (paper-based). Figure 7 shows the incident report document (hot rolling mill) to be filled out by a maintenance worker in case that an anomaly has been detected.

3.3 Outlook: Anomaly Detection in the Energy Domain

According to related surveys, integrated data analytics in the energy and resources domain promise business impact: in comparison to other industries, companies in the energy and resources industry are expected to generate, amongst others, the highest returns on big data investments. In order to make use of

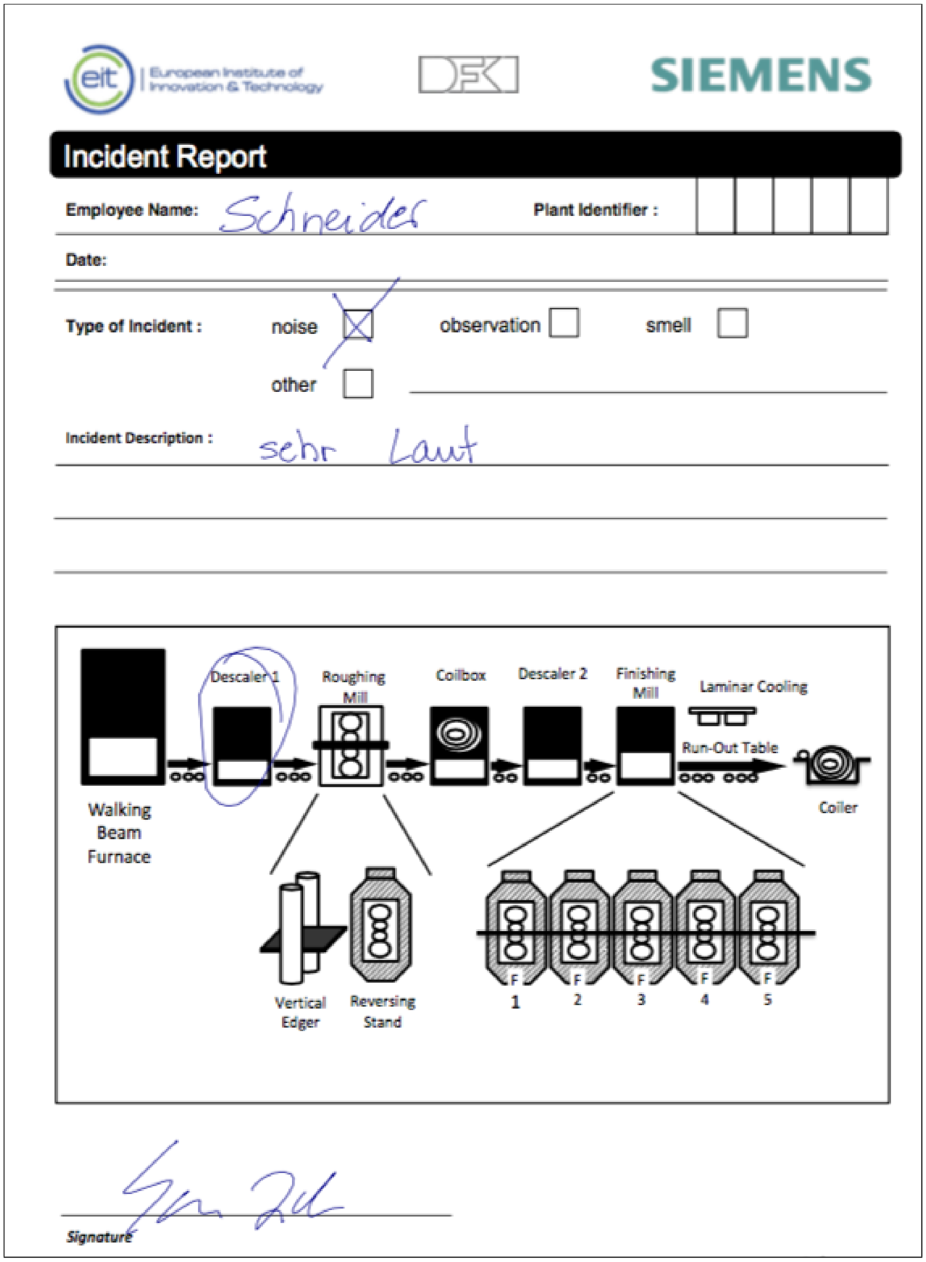

Figure 7: Incident Report Document

this wealth of data, we have two challenges to tackle:

1. seamless data integration: how to make disparate and incompatible datasets usable, interoperable and valuable across enterprises.

2. data analytics for insights into new products.

By aligning existing data discovery technologies and semantic technologies, new insights in the areas of smart meter and smart grid analytics application will be investigated and prototypically implemented. The focus will be on seamless integration of the Information Technology with the Operational Technology covering all kind of sensors from for example protection devices, via supervisory control and data acquisition (SCADA), etc.). The range of analytics applications in the integrated IT/OT world cover, for instance, voltage map generation, outage prevention, optimization applications, fuse dimensioning and asset life, asset performance management, prediction, or fault grid analysis. Potential business applications range from asset performance management, SCADA based data analysis, to outage prevention, etc. Through systematic customer evaluation processes, value propositions will be identified and implemented as prototypical and advanced functionality.

4 Acknowledgments

We thank our project partners in CPS for Smart Factories, funded by EIT Digital (EU, Horizon 2020). See http://dfki.de/smartfactories/.

References

[1] der Arbeitsdirektionen Stahl, A.E.M.: Integration von produktion und instandhaltung. Arbeitspapier, Hans-Boeckler-Stiftung (2003)

[2] Arora, S., Bhaskara, A., Ge, R., Ma, T.: Provable bounds for learning some deep representations. arXiv preprint arXiv:1310.6343 (2013)

[3] Bayer, J., Karl, M., Korhammer, D., van der Smagt, P.: Fast adaptive weight noise. arXiv preprint arXiv:1507.05331 (2015)

[4] Bayer, J., Osendorfer, C., Urban, S., van der Smagt, P.: Training neural networks with implicit variance. In: Neural Information Processing. pp. 132–139. Springer (2013)

[5] Bengio, Y.: Learning deep architectures for ai. Foundations and trends® in Machine Learning 2(1), 1–127 (2009)

[6] Bin Jabar, H.B., Sdn Bhd, S.P.: Plant maintenance strategy: Key for enhancing profitability (2003)

[7] Bleakley, K., Vert, J.P.: The group fused lasso for multiple change-point detection. arXiv preprint arXiv:1106.4199 (2011)

[8] Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? Journal of the ACM (JACM) 58(3), 11 (2011)

[9] Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Comput. Surv. 41(3), 15:1–15:58 (Jul 2009), http://doi.acm.org/10.1145/1541880.1541882

[10] Chen, T., Goodfellow, I.J., Shlens, J.: Net2net: Accelerating learning via knowledge transfer. CoRR abs/1511.05641 (2015), http://arxiv.org/abs/1511.05641

[11] Dasiopoulou, S., Giannakidou, E., Litos, G., Malasioti, P., Kompatsiaris, Y.: Knowledge-driven multimedia information extraction and ontology evolution. chap. A Survey of Semantic Image and Video Annotation Tools, pp. 196–239. Springer-Verlag, Berlin, Heidelberg (2011)

[12] Gebhard, P.: Controlling interaction with digital product memories. In: Wahlster, W. (ed.) SemProM, pp. 243–259. Cognitive Technologies, Springer Berlin Heidelberg (2013)

[13] Haupert, J., Schneider, M.: The object memory server for semantic product memories. In: Wahlster, W. (ed.) SemProM, pp. 175–189. Cognitive Technologies, Springer Berlin Heidelberg (2013)

[14] Hinton, G.E., Osindero, S., Teh, Y.W.: A fast learning algorithm for deep belief nets. Neural Computation 18, 1527–1554 (2006)

[15] Jeni, L.A., Lorincz, A., Szab, Z., Cohn, J., Kanade, T.: Spatio-temporal event classification using time-series kernel based structured sparsity. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision - ECCV 2014, Lecture Notes in Computer Science, vol. 8692, pp. 135–150. Springer International Publishing (2014)

[16] Jimenez Rezende, D., Mohamed, S., Wierstra, D.: Stochastic Backpropagation and Approximate Inference in Deep Generative Models. ArXiv e-prints (Jan 2014)

[17] Judd, S.: On the complexity of loading shallow neural networks. Journal of Complexity 4(3), 177–192 (1988)

[18] Kingma, D.P., Welling, M.: Auto-encoding variational Bayes. arXiv preprint arXiv:1312.6114 (2013)

[19] Kroener, A., Haupert, J., de Gea Fernandez, J., Steffen, R., Kleegrewe, C., Schneider, M.: Supporting interaction with digital product memories. In: Wahlster, W. (ed.) SemProM, pp. 223–242. Cognitive Technologies, Springer Berlin Heidelberg (2013)

[20] Lőrincz, A., Milacski, Z.Á., Pintér, B., Verő, A.L.: Columnar machine: Fast estimation of structured sparse codes. Biologically Inspired Cognitive Architectures 15, 19–33 (2016)

[21] Mihaly, A., Milacski, Z.A., Toser, Z., Lorincz, A., Tettamanti, T., Varga, I.: Cost and risk sensitive decision making and control for highway overtaking maneuver. In: Models and Technologies for Intelligent Transportation Systems (MT-ITS), 2015 International Conference on. pp. 87–92. IEEE (2015)

[22] Milacski, Z.Á., Ludersdorfer, M., Lőrincz, A., van der Smagt, P.: Robust detection of anomalies via sparse methods. In: Neural Information Processing. pp. 419–426. Springer (2015)

[23] Palotai, Z., Lang, M., Sarkany, A., Toser, Z., Sonntag, D., Toyama, T., Lörincz, A.: Labelmovie: Semi-supervised machine annotation tool with quality assurance and crowd-sourcing options for videos. In: 12th International Workshop on Content-Based Multimedia Indexing, CBMI 2014, Klagenfurt, Austria, June 18-20, 2014. pp. 1–4. IEEE (2014), `http://dx.doi.org/10.1109/CBMI.2014.6849850`

[24] Pistofidis, P., Emmanouilidis, C., Papadopoulos, A., Botsaris, P.: In: Proceedings of the European conference of the prognostics and health management society (2014)

[25] Schmidhuber, J.: Deep learning in neural networks: An overview. Neural Networks 61, 85–117 (2015)

[26] Schoeffmann, K., Hudelist, M.A., Huber, J.: Video interaction tools: A survey of recent work. ACM Comput. Surv. 48(1), 14:1–14:34 (Sep 2015), `http://doi.acm.org/10.1145/2808796`

[27] Sonntag, D.: Ermed - towards medical multimodal cyber-physical environments. In: Foundations of Augmented Cognition. Advancing Human Performance and Decision-Making through Adaptive Systems - 8th International Conference, AC 2014, Held as Part of HCI International 2014, Heraklion, Crete, Greece, June 22-27, 2014. Proceedings. pp. 359–370 (2014)

[28] Sonntag, D.: Tutorial 3: Intelligent user interfaces. In: Proceedings of the 2015 IEEE International Symposium on Mixed and Augmented Reality. pp. xxxii–xxxii. ISMAR '15, IEEE Computer Society, Washington, DC, USA (2015)

[29] Sonntag, D., Toyama, T.: On-Body IE: A Head-Mounted Multimodal Augmented Reality System for Learning and Recalling Faces. In: 9th International Conference on Intelligent Environments (IE). pp. 151–156 (2013)

[30] Sonntag, D., Weihrauch, C., Jacobs, O., Porta, D.: THESEUS Usability Guidelines for Use Case Applications. Technical report, DFKI and Federal Ministry of Education and Research Germany (4 2010)

[31] Sonntag, D., Zillner, S., Chakraborty, S., Lörincz, A., Strommer, E., Serafini, L.: The medical cyber-physical systems activity at EIT: A look under the hood. In: 2014 IEEE 27th International Symposium on Computer-Based Medical Systems, New York, NY, USA, May 27-29, 2014. pp. 351–356. IEEE Computer Society (2014)

[32] Tősér, Z., Lőrincz, A.: The cyber-physical system approach towards artificial general intelligence: The problem of verification. In: Artificial General Intelligence, pp. 373–383. Springer (2015)

[33] Turunen, M., Sonntag, D., Engelbrecht, K., Olsson, T., Schnelle-Walka, D., Lucero, A.: Interaction and humans in internet of things. In: Abascal, J., Barbosa, S.D.J., Fetter, M., Gross, T., Palanque, P.A., Winckler, M. (eds.) Human-Computer Interaction - INTERACT 2015 - 15th IFIP TC 13 International Conference, Bamberg, Germany, September 14-18, 2015, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 9299, pp. 633–636. Springer (2015)

[34] Šarić, M.: Libhand: A library for hand articulation (2011), `http://www.libhand.org/`, version 0.9

[35] Wahlster, W.: The semantic product memory: An interactive black box for smart objects. In: Wahlster, W. (ed.) SemProM - Foundations of Semantic Product Memories for the Internet of Things, pp. 3–21. Cognitive Technologies, Springer (2013)

[36] Zhou, A., Yu, D., Zhang, W.: A research on intelligent fault diagnosis of wind turbines based on ontology and fmeca. In: Advanced Engineering Informatics (2015)