acs plus Datengedanken Nummer 8

Approach to negotiate adequate data processing

Risk-value assessment for Cyber-Physical Systems

Case Study

acs plus
data with care

ACKNOWLEDGEMENTS

This research was partially supported by the ECSEL Joint Undertaking (JU) under the program ECSEL-Innovation Actions-2018 (ECSEL-IA) for research project CPS4EU under grant agreement No. ID-826276.

This paper documents the work done to fullfil the requirements UC1-ETH-01, UC1-ETH-02 and UC1-ETH-03.

The authors want to specially thank Alexander Mages, Aline Blankertz, Ayad Al-Ani, Bastian Halecker, Carina Mieth, Ernst von Stegmann, Ingo Sawilla, Jens Ottnad, Marit Hansen, Stefan Jähnichen, Stefan Vilsmeier, Stephan Fröhlich and the acs plus team for the fruitful and insightful discussions over the time it took to develop the concept and write this paper.

Impressum
Editor: acs plus GmbH, Rahel-Hirsch-Str. 10, 10557 Berlin
Series: acs plus Datengedanken
Number: 9
ISSN: 2941-1769
Authors: Christin Schäfer, Philipp Waack supported by Tommaso Breschi
Layout: Nadja Krombach

1. SENSITIVE INFORMATION IN CPS

1.1. PURPOSE OF THIS DOCUMENT

In this document, the three tasks UC1-ETH-01, UC1-ETH-02 and UC1-ETH-03, category ethical requirements, are discussed. The results of this analysis are presented.

1.2. ETHICAL REQUIREMENT AND CODE OF CONDUCT

Ethical requirements are concerned with the behavior of individuals, the society and its organizations, including companies. For companies, ethical requirements could influence the way the company operates as well as how their products and services are designed and delivered.

However, companies rarely refer to “ethics”. The term “code of conduct” is more commonly used in the business environment. The code of conduct reflects how a company sets its standard of behavior: What is the standard when interacting with other companies? What behavior makes the company a valuable member of society? What measures can be set to protect the environment? And especially, how should the company behave when interacting with individuals; as employees, customers or employees of other companies? In this regard, ethics and code of conduct clearly overlap.

A commonality of ethics and code of conduct is that they are only partially captured by laws. This begs the question: "why is not everything regulated through laws?" In some aspects, our society allows individuals freedom of choice. Not everything needs to be regulated. Moreover, for new developments our society must first agree on ethical standards. Only after a consensus is reached, ethical standards can be transferred into law. An example of such a new development is the so called digital transformation. One result of the digital transformation in the industrial sector is the adoption of cyber-physical systems (CPS).

Cyber-physical systems have both physical and digital components. 3D printers, CNC routers, and most machines with sensors constitute a CPS. The digital element of every CPS handles data in some way. This data might be shared and accessed by other machines or processes. From a code of conduct perspective, such sharing and accessing affect both companies and individuals: The data might contain sensitive information about individuals or companies that should not be accessible to others. This is why ethical requirements play such a big role for cyber-physical systems. Since the focus of current regulations lies on the protection of personal information, companies should take the initiative and lead by example to protect all sensitive information.

The company's code of conduct should therefore include guidelines on how CPS' handle data.
As different stakeholders have different interests and claims regarding these data, the code of conduct should also ensure that the perspectives of all stake-

holders are considered in a fair way. In many circumstances a negotiation between the stakeholders is required to really reflect what all parties consider as "fair".

The negotiations on this topic can largely vary in form and complexity. To ground the discussion, this analysis considers a real world scenario of CPS usage: a use case from the CPS4EU Horizon 2020 project: the TRUMPF Use Case Material Flow Planning and Optimization.

This use case analysis is a contribution on how to handle data ethically under a European-centric code of conduct.

Please note: The focus of this paper is the handling of data connected to CPS. The European Code of Conduct includes other topics relevant for CPS. One is the protection of the environment. All CPS must contribute to environmental protection. In the best-case scenario, the resulting products serve the environment. It must always be resource- and energy-saving, reparable, reusable and recyclable.

1.3. TRUMPF USE CASE: MATERIAL FLOW PLANNING AND OPTIMIZATION

The base for our discussion is the TRUMPF use case: "Material Flow Planning and Optimization" as described in the Grant Agreement: "All parts of the production, machine states, workers are part of a complete digital model (digital twin) of the shop floor. Through novel CPS technology interaction, this model is used for organizing, real-time controlling, forecasting as well as local and global scheduling of production processes."

The use case allows a simulation of the shop floor. It is a service used to optimize production processes. To reach that goal a digital representation of the physical shop floor is needed. This is called a static digital twin of the shop floor. The static digital twin needs to be enriched by dynamic information about processes and material characteristics.
The static and dynamic information about the shop floor are essential for the use case.

To keep this analysis as simple as possible, only two stakeholders are considered: the shop floor owner, and the service provider. The shop floor owner is the end user of the simulation services. The provider offers the simulation as a service.

Representation as data pipelines
This analysis focuses on the handling of data. Consequently, a data perspective is adopted. The shop floor simulation use case with all the information involved can be understood as a collection of data pipelines. The four data pipelines are shown in illustration 1. The generated and involved data sources are described in illustration 2.

Generation of semantic map (data pipeline 1)
The aim of the data pipeline 1 is the generation of a semantic map of the shop floor. It is created by automatically processing a 3D scan of the shop floor (data source A). This 3D information is used to infer the location of walls, doors, machines, pathways, storage areas, and obstacles. To not only locate but identify the machines, photographs of the shop floor B are also automatically processed. The information gained includes the type of machine, the producer, or the model. As output of the data pipeline 1, the semantically enhanced layout of the shop floor, named as sematic map (data source C), is obtained.

Generation of material flow information (data pipeline 2)
Data pipeline 2 is concerned with generating the dynamic information required for the simulation. In this use case, the material flow is tracked through a special Ultra-Wideband infrastructure. The tracking data

Development of the Simulation Service Toolbox
The analysis in this document focuses on the in-use phase of the simulation service. While in-use, the pipelines that create and run the simulation are fully automatized processes (illustration 1). Not a single person looks at the data during this procedure. Only the simulation outputs need to be manually reviewed. This differs from the development phase of the service. During development, several people need to validate the data. The validation can include visualizing and analyzing inputs and outputs. Therefore, the value and risk in this phase of data handling need a separate analysis. The development phase will not be the subject of this paper.

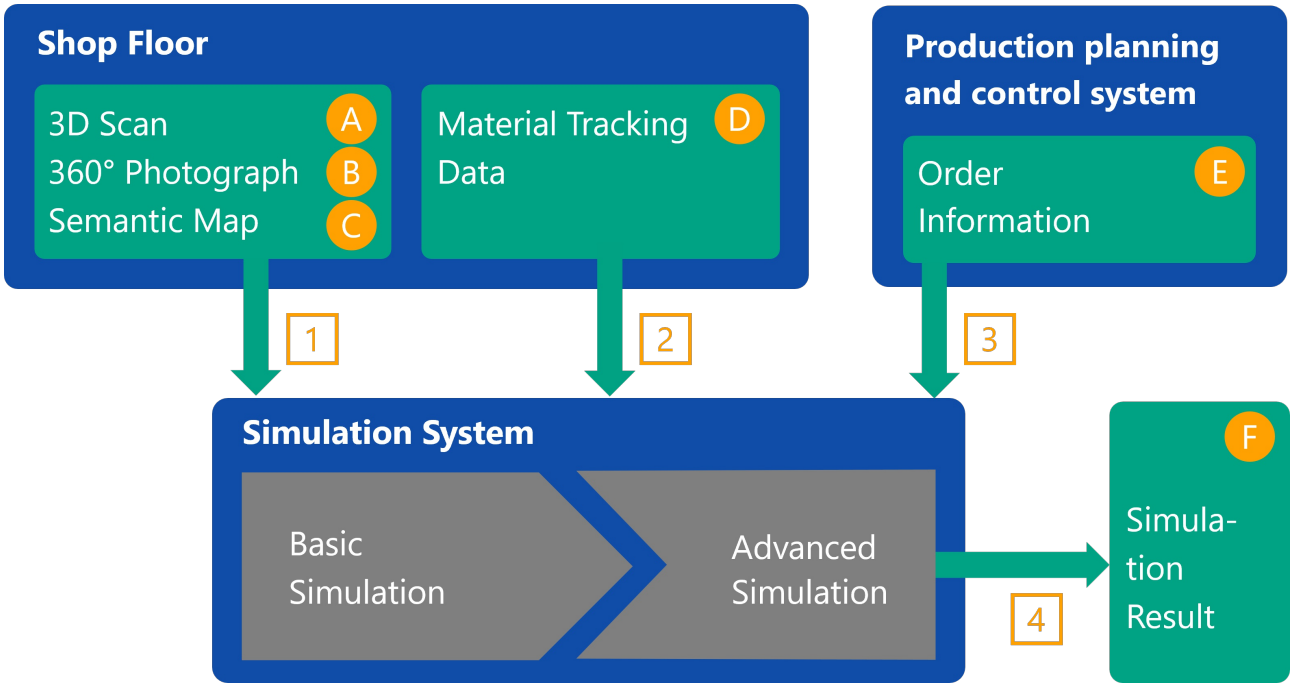

Illustration 1: The simulation service consists of four data pipelines working together to produce in the end the simulation result Ⓕ. All data sources are labelled with a letter, from Ⓐ to Ⓕ, while the data pipelines can be identified with an arabic number [1] to [4]. The first data pipeline uses a 3D scan and 360° photographs to create a digital twin of the shop floor. The second data pipeline uses material tracking information to add the dynamic component of the processes to the static digital twin. From the order data, that are extracted from a production planning and control system, detailled information about the material itself and potential production options like express can be obtained. The information gained from data pipeline [3] supports in understanding the requirements of the production processes. The output Ⓒ, Ⓓ and Ⓔ from the first three pipelines are input into the simulation framework of the service provider. The framework allows for basic as well as advanced simulation; it constitutes the data pipeline [4]. In the end, the simulation result Ⓕ is obtained.

Ⓓ is automatically collected via trackers connected to the moving pieces and raw materials. The generation of the material flow information constitutes data pipeline [2].

Order data (data pipeline [3])

For the simulation, the knowledge of specific characteristics of the material, such as the weight or dimension of certain pieces, is valuable. These properties have major influence on the simulation and the intended optimization. In this use case, it is assumed that the required information is contained in the order data Ⓔ. The order data are stored in a production planning and control system of the shop floor owner. The retrieval of this data for the purpose of the simulation is a data pipeline [3].

Simulation of shop floor (data pipeline [4])

The semantic map Ⓒ, the material flow Ⓓ, and the order data Ⓔ are all input information for the final simulation. This final processing constitutes data pipeline [4].

Different types of simulations are possible based on this information. They can be set up according to the stakeholders' needs. One question where a simulation can provide insight is finding out the quantity and type of machines in the shop floor. For this purpose, the semantic map alone can be used. More complex simulations can use multiple data sources at once. For example, to optimize the placement of the machinery the semantic map, the material flow and the order information are required. Either way, the final output is a simulation result Ⓕ.

ID	Data	Description
Ⓐ	3D scan	The scan is a 3D point cloud representation of the shop floor.
Ⓑ	360° photographs	Multiple photographs of the shop floor. They are connected to the point cloud by orientation and position.
Ⓒ	Digital twin and semantic map	Digital twin and semantic map contain description of the shop floor with pathways, storages, buffer zones, free spaces and machinery.
Ⓓ	Material tracking	Location, id and timestamp of the markers indicating stationary and moving positions of certain pieces manufactured for an order.
Ⓔ	Order information	Order specific information. It contains material and order details.
Ⓕ	Simulation results	Depending on the scenario, results contain numerical results or visualizations.

Illustration 2: The table describes the data generated for the simulation and those added from other sources. All data are generated for the use case except from data source Ⓔ. It is primarily generated for a different purpose.

1.4. LEGITIMATE INTERESTS AND CO-GENERATION

In the use case at hand, the generation of all data sources (except the order data Ⓔ) is a co-generation process. In co-generation, two or more stakeholders work together to generate the data. This implies that no stakeholder can generate data on their own. The shop floor owner can not gather data on the machines without the technology of the service provider. And likewise, the technology of the service provider is useless until a shop floor owner allows its installation and usage.
In essence, only when both parties – the shop floor owner and the service provider – work together, can data, and thus a basis for insights, be created. Neither of them could generate the result Ⓕ on their own.

From the "co" in "co-generation process", follows that both stakeholders – as generation partners – have certain rights on the data. This means that no single stakeholder is exclusively entitled to the data. Rather, the access and sharing of the information contained in the data sources should always be discussed and agreed upon between the parties. Since the stakeholders' interests in the data can vary, this discussion is crucial.

From a legal perspective, it can be be stated that the service provider has a legitimate interest in the input data Ⓐ-Ⓔ, as these data are necessary for the simulation. But the shop floor owner could also have legitimate interests beyond just Ⓕ. The intermediate data sources Ⓐ-Ⓓ, for example, can help the shop floor owner answer other business questions (e.g. the current machine utilization). In any case, it is clear that, given the co-generation, the shop floor owner has certain rights upon the data as well.

1.5. RISK PERCEPTION AND ACCEPTANCE

As it will be discussed in the next sections in more detail, generating, processing, storing, and sharing data is not free of risks. The stakeholders should always be aware of these risks and evaluate them carefully. For every use case the question needs to be answered: Does the expected value outweigh the perceived risks attached to the service?

Data sources Ⓐ-Ⓔ can contain sensitive information. There is always the risk that this sensitive information is misused. Depending on the situation at hand, the potential costs or other consequences might differ in type and significance.

The risk contained in a particular data situation is objective. Either a certain risk is present or not. If one has data stored, these data can be stolen or can be used for an unauthorized purpose.
The potential damage or loss related to a risk scenario is objective as well, even though the damage is not always precisely estimable. If personal data is stolen, a loss of reputation occurs and fines are imposed as defined in the GDPR.
On the other hand, the willingness and ability to take a risk depends on individual and subjective assessments and considerations. The risk perception varies from individual to individual. Everybody has a different risk appetite, defining the amount of risk one is willing to take.

Given a certain scenario, it is often possible to lower the inherent risk by undertaking so called risk mitigation actions. Stored personal data could be encrypted; the access to the storage could be strictly controlled via access rights. Generally the implementation of risk mitigation actions come at a cost. Therefore, whether risk mitigating actions are implemented or not depends again on the individual risk appetite.
In this upcoming analysis, some risk mitigating actions are proposed. If and to what extent the risk mitigating actions are implemented depends on the stakeholders.

1.6. SENSITIVE INFORMATION

Sensitive information contained in data involves risk. Thus, generating, processing, storing, and accessing data always involves risks. These risks differ depending on many factors; e.g. the characteristics of the data pipeline.

The evaluation of the risks, or more correctly the assessment of the expected associated harm, cleary depends on the stakeholder, and might vary significantly between them. Hence, the perspectives of both stakeholders – the shop floor owner, and the service provider – must be taken into account when their claims upon the data are negotiated.

Depending on the type of sensitive information different risks are involved. For the purposes of this document, two types of sensitive information are distinguished: Personal and business related.

[Personal information] is every information that can be traced back to an individual. A picture of a specific employee is an example. In cyber-physical systems, personal information is often not directly visible but rather hidden in technical data. In Europe, the GDPR regulates the handling of personal information.

[Business sensitive information] includes all information that are of value for the company. A patented design for a product is an example. But also economic results, business processes, trademarks, designs, or methods, belong to the business sensitive information class. So can be the number and type of CPS a competitor uses. Nationally as well as internationally, a set of rules exists for the protection of business sensitive information.

For a given data source it is not always obvious that sensitive information is present. This is particularly true in cyber-physical systems.

Sensitive information can be contained in the data in two ways: directly or indirectly. The portrait of a person is an example of data carrying direct information. Indirect sensitive information, instead, is inferred from a primary data source. It can be inferred at least in two ways: by processing the data or by linking it with external sources. For example, after processing, the path of a laser cutter could reveal the shape cut out. Extracting sensitive information via linkage is a different process. One example can be the information that the emergency button of a machine was pressed. The event time, alongside a staffing plan, can reveal the employee who pressed it.
It is difficult to identify what sensitive information is contained in a given data source. This is especially true for indirect information. Identifying sensitive information in the data is the primary challenge. However, identifying every inferable information from a data source is not possible. Furthermore, there is no way of confirming that every scenario has been considered.

1.7. RISK-VALUE TRADE-OFF

Besides risks, data offer a lot of possibilities to create value. The analyzed use case can only be successfully implemented when the data sources Ⓐ-Ⓔ are available to the service provider. Nonetheless, the expected added value of data processing does not justify sim-

ply ignoring the risks.
The expected risk and the expected value must be compared to one another. Is it worth to process the data? Does the value outweigh the risk? The answer depends on the application at hand and, as discussed in the previous section, the stakeholders' risk appetite. Since the risk profile differs from situation to situation and the risk perception and risk tolerance of the respective stakeholders varies, it is necessary to carry out a joint risk-value analysis for each constellation. There is no general assessment that applies to all parties and all situations. In particular, the willingness to implement risk-reducing measures will vary greatly between the various interest groups.

The use case specific risk-value assessment can result in different decisions on how to use and share the data. To give a few examples:

[No implementation] One option is to decide to not use the data at all neither generate, nor process, nor store, nor share the data. The risk perceived by one stakeholder may outweigh the value.

[Data minimization] The assessment might show that some of the considered data is not be essential for the use case. In this instance it is not worth generating, using, or storing these data at all. This measure implements the data minimization principle.

[Need-to-know] If the data source is essential for an application, the need-to-know principle should be applied. The parties must agree on which data points are actually necessary for which stakeholder in which processing step of the use case. The necessary agreements define, for example, the level of detail or the frequency of data delivery.

1.8. DATA AGREEMENTS AS PART OF SERVICE CONTRACT

In many CPS scenarios, there is a B2B relationship between the stakeholders. This definitely applies to the analyzed TRUMPF use case. It means that there is usually a contract between the stakeholders. Without a contractual arrangement, the service provider will not install its technology at the shop floor. Similarly, the shop floor owner will not grant the service provider access to his or her premises. It is advisable to contractually fix the data agreements as part of such a service contract. This includes the data usage agreement, i.e. the definition of those questions for which the data may be analyzed, and the measures for data minimization or risk-mitigation.

It can be assumed that a few standard clauses will emerge for the implementation of a certain CPS service or product.

The service provider should limit his or her claims to the data. Ideally, to the minimum demonstrably necessary to implement the use case. Evidence of strict data minimization must be provided and becomes part of the contract. In this situation, the options for the shop floor owner are reduced to agreeing to the contract or forgoing the service. The contract itself is standardized.

A standardized extension of the contract can state whether the contractual partner agrees to other secondary uses of the data. Such a secondary benefit could be the further development of the service by the service provider. If a stakeholder agrees to an extended use of the data, the other partner must also provide a compensation in return.

In this document, the risk-value assessment is – for illustrative purposes – carried out for the TRUMPF Use Case. The perspectives of both stakeholders – the shop floor owner and the service provider – are considered. The risk-value assessment is structured around the four data pipelines (1, 2, 3, 4) in illustration 1. The pipelines and their respective data sources will be analyzed separately.

2. RISK AND VALUE ASSESSMENT

2.1. DATA PIPELINE 1: GENERATION OF SEMANTIC MAP

The data pipeline 1 as described in section 1.3. is the first in focus. Data pipeline 1 generates the semantic map C by processing the 3D scan A and the photographs B of the shop floor. In the upcoming section, a risk-value assessment of the data sources A, B and C is performed. This assessment starts with analyzing the data for presence of sensitive information. Given the fact that the data are necessary for the intended shop floor simulation, they possess a significant value. After the potential risks are identified, the risk-value assessment is performed. Where applicable, recommendations for risk mitigating actions are given.

2.1.1. ANALYSIS OF BUSINESS SENSITIVE INFORMATION

As described in section 1.6., business sensitive information are of value for the company and should be safeguarded to avoid risky exposure of the information.

Whether and to what degree there is a risk depends on whether and to what degree data sources A, B and C contain business sensitive information. The following chapters analyze this question individually for each of the data sources mentioned.

2.1.1.1. 3D SCAN OF THE SHOP FLOOR

In this use case implementation, the 3D scan A is generated in multiple steps. First, different sensors measure the shop floor simultaneously; a gyroscope measures acceleration and orientation of the measurement device, multiple cameras take pictures, and a laser scanner measures distances in the shop floor. Second, all sensor data are merged in a complex sensor fusion. The result of the sensor fusion is a point cloud that forms the basis for the following data processing.

One exemplary point cloud in this use case consists of 81.866.667 data entries. The stored file has a size of 6.8 Gigabyte. Each point of the point cloud is represented via one line of the file. Each line consists of six variables: the location of the point – x, y, z coordinates – and the color of the point – R, G, B values. The coordinates are decimal numbers in a coordinate system. A coordinate can be of negative or positive value. The color data are integers in the RGB color space ranging from 0 to 255. An example of a 3D scan data structure is given in illustration 3.

The description of the 3D scan A and illustration 3 make it clear that no human can draw a conclusion by visually inspecting the raw data. It is impossible for a human to recognize a wall, a machine, or a door in this raw data. And it is just as impossible for a person to directly read the business sensitive information associated with this data. Therefore, it is not only efficient, but fundamentally necessary that the 3D scans are evaluated fully automatically using specially developed algorithms. 3D scans only release their information "indirectly" – as defined in section 1.6. – through

Location			Color		
X	Y	Z	R	G	B
-22.98	31.75	3.64	179	156	138
5.46	5.08	-0.82	153	170	166
-34.40	34.32	3.31	183	169	158
-4.05	5.45	2.16	110	96	82
2.72	-0.74	2.31	191	193	197
...	...	...	...	...	...

Illustration 3: Extract of the raw data from a 3D scan Ⓐ. For every location identified via the coordinates x,y and z, that the scan reached, the observed color is given in the RGB coding. With this raw data, a human cannot infer any part of the layout by simply looking at it.

processing. This applies to the primarily required information on the layout of the shop floor, as well as to the business sensitive information inherent in this layout.

Business sensitive information

[Direct] From the above explanations, it can be seen that no business sensitive information is contained directly in the raw data of the 3D scans Ⓐ.

[Indirect] However, business sensitive information can be indirectly inferred by processing the scan. The resulting layout of the shop floor itself is business-sensitive information. For example, the layout can be used by a competitor to gather intel on the equipment of the shop floor and on production processes.

Value

[Essential] The 3D scan Ⓐ is essential for generating the semantic map that is itself necessary for running the simulation service. Obtaining the layout for the semantic map is the only purpose of generating the 3D data set of the shop floor. Thus, in this use case, business sensitive information is essential.

[Data minimization] The 3D scan Ⓐ contains the minimum information required to derive the layout of the shop floor. A reduction of the data, either with regard to the recorded variables, or with regard to the scope of observations, is not possible. If the shop floor is to be simulated, the data must be collected in the presented form.

Risk

- At the moment the 3D scan Ⓐ is created, the risk arises for the shop floor owner that this data and the inherent business sensitive information will be used in an undesirable way.
- As no information about the service provider is contained in the 3D scan Ⓐ, the service provider is not exposed to risks.

Use case inherent risk mitigation

In this use case, the process of inferring the layout from the 3D scan Ⓐ of the shop floor is automatic. No person manually handles the scan data.

[Who] This reduces the involved parties, and therefore the number of potential criminal actors.

[How] An automated process limits the opportunities to use the data in an undesirable way.

Given the characteristics of the use case implementation, the efforts required to steal the data are higher compared to a process that involves interactions by humans or additional other processes.

Aspects for negotiation

Generally, the risk that comes with inferring the layout information must be taken by the shop floor owner. Otherwise, the service cannot be executed. If this risk is accepted, there are the following aspects for negotiation:

[Data usage] Should the 3D scan Ⓐ be used only for the derivation of the layout or could other usages be considered as well? Is the service provider allowed to use the data for improving the algorithms used for the derivation? Can the shop floor owner give access to the data to a competitor of the service provider?

[Storage] Should the 3D scan be stored? Already prior to processing or afterwards? Where?

[Access] If the data is stored, who has access?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the service provider is only allowed to use the 3D scan Ⓐ to derive the layout for the semantic map.
- the data are not stored at all and deleted directly after processing.
- only the process to derive the layout has ever access to the data.

2.1.1.2. 360° PHOTOGRAPHS OF THE SHOP FLOOR

In this use case the 360° photographs Ⓑ are used to identify the machines in the shop floor. While from the layout the location of a machine can be derived, with the photographs it becomes possible to learn specific details about the machines. These details encompass the general type of the machine, the specific model, or the producer. The photographs are accompanied by information on the position they are taken from in the shop floor, and their orientation. This additional information allows to combine the 3D scan Ⓐ and the photographs automatically.

To analyze one shop floor, more than a hundred 360° photographs are taken, the exact number depends on the size and the structure of the shop floor. With the equipment used in this use case, such a photograph has a size of about 6 MB.
Each file consists of 33.554.432 pixel, stored in 3 large table with 4096 rows and 8192 columns; one for every dimension of the RGB color space. As already shown in section 2.1.1.1, the RGB information is encoded by values between 0 and 255; one value reflecting the red, green, and blue portion of the pixel. illustration 4 shows an example of such a 360° photograph data structure.

To derive the information required to identify the machines, in this use case, the photographs are processed automatically. The photographs are processed only as a series of RGB numerical values. They are not analyzed in a visual form.

As with the 3D scans Ⓐ, the 360° photographs Ⓑ are also illegible for humans in their raw digital form. No human can find out the type or model of a machine from a data format as shown in illustration 4. Again, it is not only efficient, but fundamentally necessary that the 360° photographs are evaluated fully automatically using specialized algorithms. The information contained in the photographs is only obtained indirectly via processing. This applies to the primarily required information on type, model, and producer of the machine, as well as to the business sensitive information inherent in these characteristics.

Business sensitive information

[Direct] The numeric RGB values, without processing and visualizing the data, do not contain any direct business sensitive information.

[Indirect] But there is indirect business sensitive information. In the use case implementation, the data are processed automatically to identify the machines. The type and model of the machines, the producer, together with other machine specifications, are business sensitive information. As a special form of processing the data, a visualization – what is normally referred to as a photograph – would reveal this business sensitive information as well. But in this use case, the data are not visualized.

R	1	2	3	4	...
1	45	42	44	43	...
2	45	46	45	44	...
3	46	44	47	56	...
4	44	53	40	34	...
5	99	43	36	39	...
...	...	...	...	...	...

G	1	2	3	4	...
1	12	9	11	10	...
2	12	13	12	11	...
3	13	14	3	1	...
4	12	18	0	0	...
5	83	4	0	0	...
...	...	...	...	...	...

B	1	2	3	4	...
1	58	55	57	56	...
2	58	59	58	57	...
3	59	57	63	73	...
4	56	61	44	38	...
5	99	40	41	48	...
...	...	...	...	...	...

Illustration 4: Extract from a 360° photograph Ⓑ showing the values for red, green and blue for every pixel the photograph consists of. With this raw data, a human cannot infer the type of a machine by simply looking at it.

Value

[Essential] Similar to the 3D scan Ⓐ in subsection 2.1.1.1, the 360° photographs Ⓑ are essential for running the simulation service.

[Data minimization] The 360° photographs Ⓑ are the minimal required set of information that allows the identification of the machines. A reduction of the raw data is not possible. If the shop floor is to be simulated, the photographs must be taken.

The possibilities of inference on the 360° photographs Ⓑ are very extensive. In addition to the derivations required for the simulation, such as the type of machine, inferences on aspects that have no relevance to the simulation are also possible. For example, the age or the condition of the machine could be revealed. Knowing the age and condition of the machine allows drawing conclusions about the company's financial health and competitiveness.

Derivations from the 360° photographs are required for two reasons:

[Required] The information is required in exactly this form for the simulation. This is the case, for example, for the type of machine.

[Support role] The information itself is not required for simulation but allows other required information to be identified. This "support" role is played, for example, by the information about the specific model of a machine and the manufacturer. With their help, For the simulation, it is important to know the type of a machine. Whether a machine is a laser cutting machine, or a bending machine, directly influences the simulation.

The situation is different when it comes to information about the model or the manufacturer of a machine. This information initially has no added value for the simulation itself. Instead, it plays a role of support. For more in-depth simulations and optimizations, it is important to know the performance characteristics of the machines. For example, the throughput or processing time for a machine type can vary from model to model. Some machines can process sheet metal up to a thickness of 15 cm, others up to 30 cm.

These performance characteristics cannot be directly inferred from the photographs. But knowing the model and/or manufacturer of the machine allows to look up this performance characteristics from publicly available data sources.

In order not to overload this analysis, the deduction of performance characteristics from model and/or manufacturer information will not be discussed further. This inference is considered an automatic step in the data pipeline [1] performed by the service provider.

it is possible to identify performance characteristics required for the simulation.

Risk

- At the moment the 360° photographs Ⓑ are created, the risk arises for the shop floor owner that this data, and the inherent business sensitive information, will be used in an undesirable way.
- As no information about the service provider is revealed from the 360° photographs Ⓑ, the service provider is not exposed to risks.

Use case inherent risk mitigation

- As described above, identifying machines is also possible when processing the data into a visual form: the "photograph", or image. On an image, every human can see the machine's type, model or producer. In this use case implementation, the raw data are not processed as images. Therefore, the sensitive information is not exposed via visualization.
- The processing of RGB data is done automatically. No human is in the loop.
  [Who] This reduces the involved parties and therefore the number of potential criminal actors.
  [How] An automated process limits the opportunities to use the data in an undesirable way.

Aspects for negotiation

If the shop floor owner wants the simulation, thus perceives the value higher than the potential harm associated to the risk, the data must be generated. But it can still be negotiated what data exactly is generated and used and who has access to what data. This leads to the following negotiable aspects:

[Data usage] Is the service provider allowed to use the 360° photographs Ⓑ for any other purpose than the simulation of the shop floor?

[Inference] What information should be inferred from the 360° photographs Ⓑ?

[Information look up] What performance characteristics are looked up from external sources?

[Providing of Output] Should the support role information be provided for the next processing step?

[Storage] Should the 360° photographs Ⓑ and/or the inferred information and/or the performance characteristics be stored? If so, all of the information including the support variables?

[Access] Who should have access to the photographs and the inferred information? Should the 360° photographs Ⓑ and/or the inferred information and/or the performance characteristics be accessible to the shop floor owner?

[Deletion] Should the 360° photographs Ⓑ be erased after the processing is completed?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the service provider is only allowed to use the 360° photographs Ⓑ to derive the layout for the semantic map.
- only information relevant either directly for the simulation or in a support role are inferred from the 360° photographs.
- only performance characteristics relevant for the simulation are looked up.
- only information relevant for the simulation is passed on.
- no data is stored.
- no access to the data is granted other than to the process extracting the required information from the 360° photographs Ⓑ.
- all information is erased directly after the processing is completed.

2.1.1.3. SEMANTIC MAP

The semantic map Ⓒ results after the processing of data sources Ⓐ and Ⓑ by data fusion.

With the 3D scan Ⓐ and the 360° photographs Ⓑ, a distinction is made between the digital raw data, and a potential visualization. An example of such a visualization is given in the illustrations 6, and 7, in the following section 2.1.2. An analogous distinction between digital raw data and visualization also makes sense for the semantic map Ⓒ.

A visual representation of the shop floor is not required for the actual simulation. Rather, the digital twin, i.e. the digital description of the shop floor and its properties, is required. The data processing of the 3D scan Ⓐ, the 360° photographs Ⓑ, and the data fusion of the results, generate exactly this digital twin of the shop floor.

A visualization of the digital twin that people can understand is the semantic map. A semantic map generated as part of this use case can be found in illus-

tration 5, which follows later.
Since the simulation itself does not need the semantic map, the question arises to what extent such a semantic map should be generated from the digital twin. The answer to this question depends on the wishes of the shop floor owner. It is undeniable that such a semantic map has great added value and benefits.

In the following discussion, a distinction is made between the digital twin and the semantic map, both of which are denoted by C.

In order to enable the simulation intended in this use case, the digital twin C, i.e. the database created through the data fusion of the results of processing the 3D scans A, and the 360° photographs B, contains at least information about the location of the machines, their type, and performance characteristics, as well as information on the overall layout of the shop floor.
Whether the digital twin C also contains other information that is not relevant for the simulation depends on the agreements made for the processing of the 3D scans A, and in particular the 360° photographs B. If the recommendations for risk minimization have been implemented, the digital twin only includes the information required for the simulation. Otherwise, information about the model, the manufacturer, or the condition of a machine might be available as part of the digital twin C.

As for data source A and B, and therefore also for the digital twin C, a human observer cannot do anything with the information by simply looking at it. An inference of business sensitive information is only possible by means of automated processing.
The situation is completely different with the visualized semantic map: the viewer immediately realizes facts that are sensitive from the perspective of the shop floor owner.

Business Sensitive Information

[Direct] As raw data, the digital twin C does not contain direct business sensitive information that humans can interpret.

[Indirect] To obtain the semantic map C from of the digital twin C, processing is required. The semantic map C contains business sensitive information. Also, by using the digital twin C as input for the simulation, business sensitive information is revealed.

Value

[Essential] The digital twin C is essential to run the simulation. Hence, in this use case, business sensitive information is essential.

[Data minimization] Whether the digital twin C contains only information relevant for the simulation, or other data as well, depends on the agreements made for the processing of the 3D scans A, and in particular the 360° photographs B.

Risk

- As soon as the digital twin C is generated, the risk arises for the shop floor owner that this data and the inherent business sensitive information will be used in an undesirable way.
- As no information about the service provider is contained in the digital twin C, the service provider is not exposed to risks.

Sharing the digital twin C with the service provider leads to the exposure of business sensitive information. This holds true even with the smallest possible set of information, as this still includes the shop floor layout and the machines with type, location, and performance information.

Machines' performance characteristics are also sensitive information: knowing a machine's pieces per hour capacity can reveal the maximum yearly production. Linking this information with prices, assessments on the company's revenue can be made. Moreover, having access to specific performance characteristics can lead to the identification of single machines. For example, there might be only one laser cutter model with a maximum speed of 130m/min on the market. Knowing this single performance characteristic could therefore help to uncover the machines' model, manufacturer, or acquisition cost.

Use case inherent risk mitigation

In this use case, the data fusion to obtain the digital twin C is done automatically without a human in the loop.

[Who] This reduces the involved parties, that is the number of potential criminal actors.

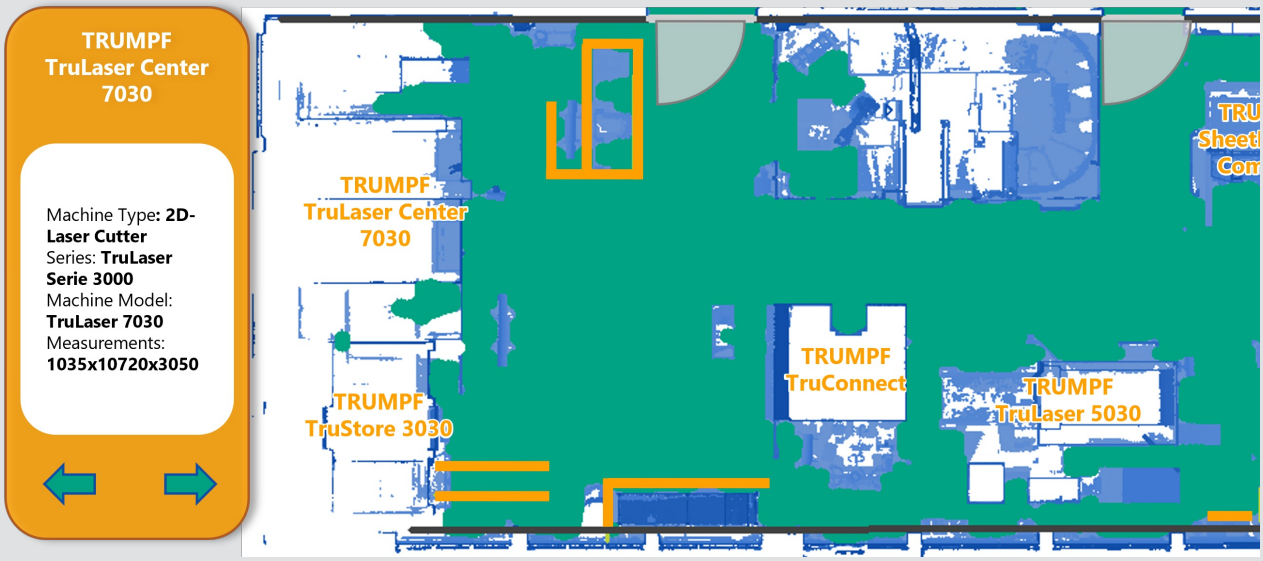

Illustration 5: Zoom in a semantic map of a shop floor: The map shows the overall layout of the facility with machines, their type, manufacturer and model together with additional performance characteristics.

From one digital twin of this use case, a semantic map is generated for illustration purposes.
From a technical perspective, the semantic map is an image. In the given example, the size of the image is 1138 x 2087 x 3. Similarly to the 360° photographs displayed in illustration 4, there is one layer of information for red, green, and blue.
From a visual perspective, the semantic map is on the one hand a map; with the layout of the shop floor, the location of machines, walls, doors, floor markings, free space, and potential obstacles clearly identifiable. On the other hand, the visualization is semantically enriched: further information about the machines, their type, model, and manufacturer are given, and there is the possibility to display even the performance characteristics.

[How] An automated process limits the opportunities to use the data in an undesirable way.

Aspects for negotiation

To run the simulation of the shop floor, parts of the information from the digital twin must be shared. The aspects to negotiate are:

[Data usage] It must be negotiated whether the service provider may use the data for purposes other than simulating the shop floor.

[Data minimization] The subset of information from the digital twin C used for the simulation must be agreed on.

[Data minimization] If performance characteristics are delivered to the simulation, the granularity of the information should be agreed on. Instead of using exact values for the characteristics, a performance class could be used. This way the ability to infer the machine model from the characteristics is limited, while the loss in quality for the simulation is small.

[Semantic map] Should a semantic map C be created, which visualizes the information from the digital twin?

[Semantic map] If a semantic map C is produced, it needs to be agreed on what information is shown. While not used in the simulation, support role information like model and manufacturer could be displayed.

[Storage] Should the digital twin C be stored? Should the subset of information sent to the simulation be stored? Should the semantic map C be stored?

[Access] Who gets access to the digital twin Ⓒ, to the semantic map Ⓒ, and to the subset of information from the digital twin used within the simulation?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk
The risk inherent to the use case can be minimized if

- the digital twin consists only of information required for the simulation.
- performance characteristics are mapped to performance classes and only these classes are sent to the simulation.
- the service provider is allowed to use the data only for the simulation.
- no data is stored.
- no semantic map is produced.
- if it is decided to create a semantic map, it should be under the full control of the shop floor owner.

2.1.1.4. SUMMARY DATA PIPELINE [1]: SEMANTIC MAP

In data pipeline [1], business sensitive information is present in all data sources considered: the 3D scan Ⓐ, the 360° photographs Ⓑ, and the digital twin Ⓒ. The 3D scan Ⓐ, and the 360° photographs Ⓑ, are necessary to generate the digital twin Ⓒ which is necessary for the simulation. The only risk-free option would be not to run the simulation at all.
If the simulation is deemed valuable (enough), the shop floor owner must take the risks attached to sharing sensitive business information about the shop floor with the service provider. The risks can be minimized following the recommendations given in the previous subsection.

The semantic map Ⓒ is a visualization of the information contained in the digital twin Ⓒ. Similarly to the 3D scan Ⓐ and the 360° photographs Ⓑ, a visualization of the raw data is not required for the simulation, as the processing is done automatically. It can be assumed that such a semantic map Ⓒ is of interest for the shop floor owner. If this is the case, the semantic map Ⓒ should be treated as an additional product, and the terms and conditions for sharing the information should be negotiated.

2.1.2. ANALYSIS OF PERSONAL INFORMATION

In section 2.1.1, the risk-value assessment of data pipeline [1] for business sensitive information was conducted. This section performs the risk-value assessment for personal information.

For the simulation of the shop floor, having the goal to optimize production processes, personal information is neither required nor wanted at all. Personal information adds no value to the simulation.
Even though it is not needed, personal information could still be contained in data sources Ⓐ and Ⓑ, and potentially inherited by Ⓒ. It is important to be aware of which exact personal information is embedded in the data.
With business sensitive information, the risk of misuse concerns the shop floor owner and the service provider only. With personal information, instead, the risk of misuse can also impact both stakeholders' employees. One potential misuse of personal information in this context could be the monitoring of employees for the purpose of performance assessments and behavior control.
In Europe, where the GDPR regulates the protection of personal data, companies must handle personal information with care. Within the value and risk assessment process, the interests of the employees are represented by their respective employer.

The analysis of personal information for data pipeline [1] will be carried out in the following subsection.

2.1.2.1. 3D SCAN OF THE SHOP FLOOR

The 3D scan Ⓐ is used to infer the layout of the shop floor. It is described in detail in subsection 2.1.1.1, and the data structure is shown in illustration 3. As explained there, the data is processed automatically without human intervention.

Analogously to the findings for business sensitive information, it also applies to personal information that no knowledge can be derived directly from the 3D scan raw data Ⓐ by visual inspection. This is only possible after certain processing is applied.

Personal information
[Direct] The raw data does not contain direct personal information.

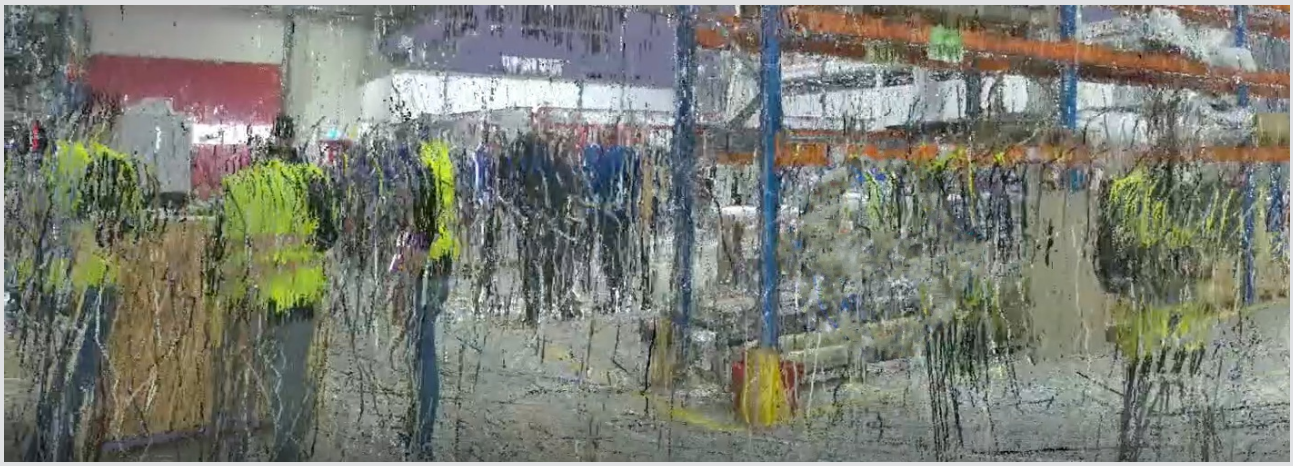

Illustration 6: The image shows an example visualization of a 3D scan Ⓐ. There are people recognizable.

The 3D scan Ⓐ can be further processed into a visualization. If a person was present at the moment of the 3D scan the visualization reveals people's silhouettes; but only silhouettes. It is important to note that from these silhouettes it is normally impossible to identify a certain person. The silhouettes indicate that a person is in the frame, but who that person is remains unknown.

The difference between recognizable and identifiable can be seen in the illustration above: There, the silhouettes of groups of workers can be seen. One can recognize they are people, but not identify any specific employee. To identify a specific employee from such a 3D scan, additional specific knowledge about him or her is needed. For example, a staffing plan could be cross-referenced with the 3D scan.

[Indirect] There can be personal information contained indirectly. If employees are present while scanning, there will be coordinates representing them in the 3D scan Ⓐ.

Value

As already mentioned in the introduction of section 2.1.2, personal information adds no value to this use case. Neither the shop floor owner nor the service provider is interested in this personal information. Hence, there is no value in having any personal information in the 3D scan Ⓐ.

Not only is personal information irrelevant for the simulation, but it could even be detrimental. If people stand in front of the machines during the scanning operations, the automatic identification process can encounter issues. The algorithm is trained to identify machines only. This means that human silhouettes can contaminate the data, resulting in artifacts as can be seen in illustration 6. Consequently, the automatic identification process could even fail.

Risk

- No risk arises for the shop floor owner or the service provider from the existence of personal information in the 3D scans Ⓐ; assuming that no general data protection regulations are violated.
- The risk of being identified via the 3D scan Ⓐ exists for the employees of the shop floor owner.

The potential impact and damage to the employee in the event of identification are rather small. The only information obtained would be that a certain employee was in a certain location in the shop floor at a given time.

Use case inherent risk mitigation

- The 3D scan Ⓐ is processed automatically without a human in the analytical loop.
  [Who] This reduces the involved parties and therefore the number of potential criminal actors.
  [How] An automated process limits the opportunities to use the data in an undesirable way.

- Moreover, the 3D scan Ⓐ is not visualized, and from the raw data humans cannot infer any meaningful information.
- The 3D scan Ⓐ itself is only taken once, and not on a continuous basis. The information potentially gained from one scan about a single employee is too limited to be used in any type of performance benchmarking.

Since the potential to harm the employee is relatively low, complex or costly measures to reduce the presence of personal information in the data are neither necessary, nor expected.

Aspects of negotiation

[Process] When should the 3D scan Ⓐ be taken? Is it possible to empty the shop floor at this point in time? The scan could be conducted at night or during lunch breaks. This would help to prevent to collect personal information and potential artifacts caused by people being present.

[Data usage] It must be negotiated whether the service provider may use the data for purposes other than simulating the shop floor.

[Storage] Should the 3D scan Ⓐ be stored? Where?

[Access] Who should have access to the 3D scan Ⓐ?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the shop floor is empty at the time of scanning;
- in the event that it is not possible to clear the shop floor, the employees have been informed and admission is legally secured via the employment contract or a company agreement.
- the 3D scan Ⓐ will be deleted after processing.
- only the relevant processing for the simulation has access to the 3D scan Ⓐ.
- the service provider uses the data only to create the layout required for the semantic map.

2.1.2.2. 360° PHOTOGRAPHS OF THE SHOP FLOOR

Along with the 3D scan Ⓐ, 360° photographs Ⓑ of the shop floor are taken. The technical details and the data structure are described in subsection 2.1.1.2, and displayed in illustration 4. As clearly pointed out, personal information is not of interest for this use case. Therefore, personal information is not actively collected, but may be included in other data sources as a kind of by-product or artifact. Exactly this case, personal information as an artifact, can arise for the 360° photographs Ⓑ: If people are present when the photographs are taken, information about them will be part of the digital representation of the 360° photograph.

Personal Information

[Direct] The numeric RGB values, without processing and visualizing the data, do not contain any direct personal information.

[Indirect] But there is indirect personal information that can be revealed by processing. As a special form of processing the data, a visualization in the form of what is normally referred to as the photograph would clearly show all people that had been present at the moment the picture was taken. An example is given in illustration 7. In turn, the possibility to identify a single person makes the photograph data inherently riskier than the scan data for the employees of the shop floor owner.

It is important to note that the implementation of the use case deliberately refrains from making the digital data of the photograph available in a visualized form. The digital automated process does not require such visualizations.

Value

No personal information from the 360° photograph Ⓑ is required for the use case.

Risk

- There arises no risk for the shop floor owner or the service provider from the existence of personal information in the 360° photograph Ⓑ; assuming that no general data protection regulations are violated.
- There is a risk for the employees of the shop floor owner to be identified from a visualization of the 360° photographs Ⓑ.

If an employee is identified from a 360° photograph Ⓑ, the potential damage for the employee is rather small. The only information obtained would be that a specific employee was in a certain location in the shop floor at a given time.

Illustration 7: An image generated from a 360° photograph Ⓑ of a shop floor with two people clearly visible, but intentially blurred in this example. A visualization like this one of the 360° photographs Ⓑ is not used in the process to run the simulation.

Use case inherent risk mitigation

- The processing of data source Ⓑ is done fully automatized. No manual analysis is carried out.
  [Who] This reduces the involved parties and therefore the number of potential criminal actors.
  [How] An automated process limits the opportunities to use the data in an undesirable way.
- The 360° photographs Ⓑ are processed in a non-visual form.
- The 360° photographs Ⓑ are taken only at a certain point in time. This makes their usage for benchmarking purposes rather unlikely.

Since the potential to harm the employee is relatively low, complex or costly measures to reduce the presence of personal information in the data are neither necessary, nor expected. The inclusion of personal information in the data should still be avoided if possible.

Aspects of negotiation

[Process] When should the 360° photographs Ⓑ be taken? Is it possible to empty the shop floor at this point in time?

[Data usage] It must be negotiated whether the service provider may use the data for purposes other than simulating the shop floor.

[Storage] Should the 360° photographs Ⓑ be stored? Where?

[Access] Who should have access to the 360° photographs Ⓑ?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the shop floor is empty at the time of scanning; In the event that it is not possible to clear the shop floor, the employees have been informed and admission is legally secured via the employment contract or a company agreement,
- the 360° photographs Ⓑ will be deleted after processing.
- only the relevant processing for the simulation has access to the 360° photographs Ⓑ.
- the service provider uses the data only to identify machines to enrich the semantic map.

In this particular use case, a special device is used to scan the hall and take photographs. The device used requires an operator. The operator might be photographed while scanning and photographing the shop floor. To avoid photographing the operator, the device used can be mounted on a remote carriage.

2.1.2.3. SEMANTIC MAP

Data sources Ⓐ and Ⓑ are automatically processed to generate the digital twin Ⓒ. The semantic map Ⓒ is a visualization of (parts of) the digital twin Ⓒ.

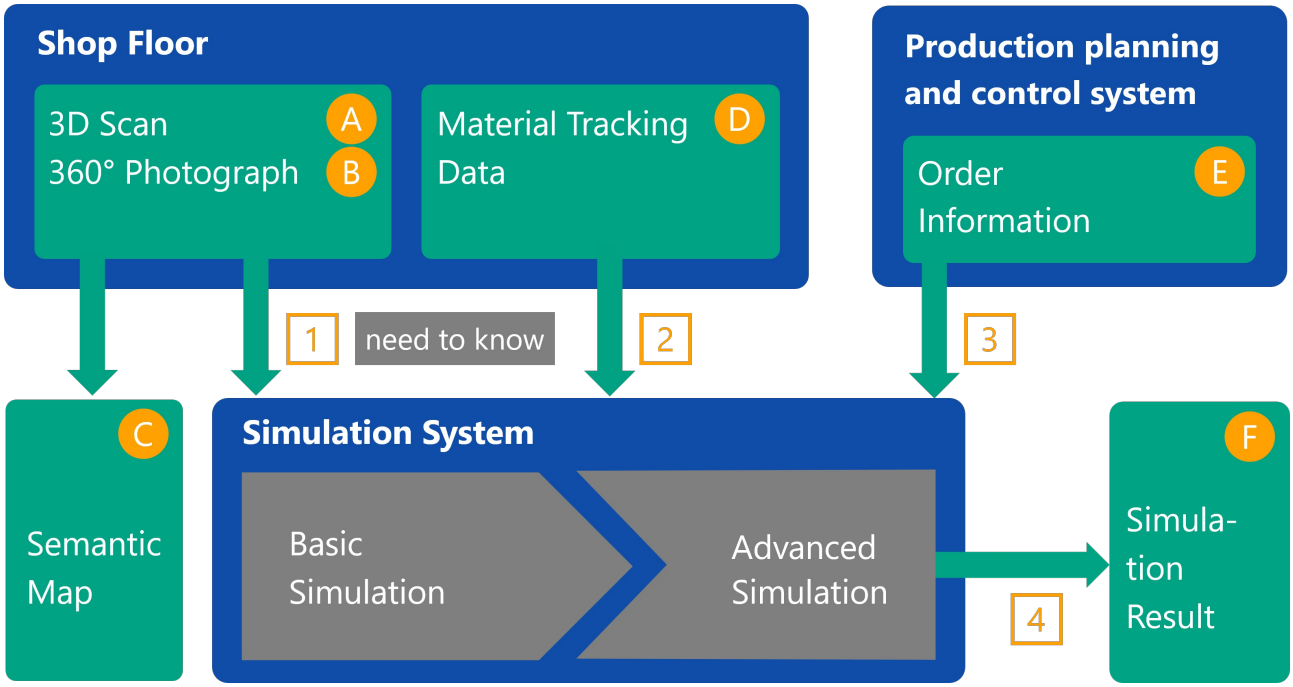

Illustration 8: Update of the data pipeline illustration 1. Only need-to-know information should be provided for the simulation and the semantic map should only be accessible for the shop floor owner.

Both the digital twin Ⓒ and the semantic map Ⓒ do not contain any personal information. The processing of the 3D scan Ⓐ and the 360° photographs Ⓑ eliminates personal information, if it had been present in the first place.

Personal information

[Direct] No direct personal information is contained in the digital twin Ⓒ.

[Indirect] It is hard to imagine a scenario where processing the digital twin and/or linking the information with another data source might reveal personal information.

As no personal information is contained in the digital twin Ⓒ a risk-value assessment is not required.

2.1.2.4. CONCLUSION

Personal information can be contained in both, the 3D scan Ⓐ and the 360° photographs Ⓑ. After the data fusion to generate the digital twin Ⓒ, the personal information is eliminated. For the simulation no personal information is required nor wanted.
It is possible to avoid the presence of personal information by taking the 3D scan Ⓐ and the 360° photographs Ⓑ when no employee is present in the shop floor. Otherwise, risk-mitigating measures should be taken as described in the subsections 2.1.2.1 and 2.1.2.2.

Given the results from the discussion of the data pipeline [1], first for business sensitive and second for personal information, the illustration 1 is updated. Illustration 8 shows one major change: the semantic map that reveals business sensitive information to every observer but is not necessary for the simulation is taken outside of the simulation process as a standalone artifact. The terms and conditions regarding the semantic map need to be negotiated. It is recommended, that the shop floor owner has full control over it.
Within the remaining process, only information essential and relevant should be provided to the simulation.

2.2. DATA PIPELINE 2: MATERIAL FLOW

The data pipeline 2 generates material tracking data allowing to simulate the material flow. In this use case the material tracking data Ⓓ is generated by placing a tracker on a trolley carrying materials. This way, the coordinates of the trolley at certain points in time are collected. In this section the resulting tracking data are analyzed for presence of sensitive information.

2.2.1. ANALYSIS OF BUSINESS SENSITIVE INFORMATION – MATERIAL TRACKING DATA

The material tracking data are used to infer the paths of the trolleys and thus, the location and path of the materials. The material traces are business sensitive information.

To be able to distinguish between trackers, every tracker has a unique ID, that is part of the data set. Beside the ID for each material trace, the timestamp of the measurement and the x, y and z coordinates of the location are given as illustrated in 9. An exemplary sample rate of 5 measurements per minute leads to 18000 observations per hour per material trace.

The description of the material tracking data and the example in illustration 9 make it clear, that except the total number of tracker in use, more substantial information like the material traces can not be inferred by any human. This knowledge can only be gathered by automatically processing the information.

Business sensitive information

[Direct] Business sensitive information is directly recognizable from the material tracking data Ⓓ, as the number of trackers in use is obvious.

[Indirect] Business sensitive information is also contained indirectly, as the material traces itself can be revealed after processing the data.

Both information allow for example conclusions about the order situation of the company as well as the size and capacity of the production facility. The number of active trolleys can give hints on the volume of orders. Few trolleys moving could mean the company is not receiving many orders.

Value

[Essential] To obtain the traces for the material, the material tracking data Ⓓ are essential. Without the traces certain optimizations cannot be performed with the simulation. Thus, in this use case, business sensitive information is essential.

[Data minimization] The data structure of the material tracking data Ⓓ cannot be reduced without losing the ability to infer the traces. ID, timestamp

ID	timestamp	x	y	z
1	2021-11-02 12:55:13	38.4	345.2	15
1	2021-11-02 12:55:14	39.4	342.2	15
1	2021-11-02 12:55:15	40.5	341.6	15
2	2021-11-02 12:55:15	384.5	32.1	15
1	2021-11-02 12:55:16	40.5	341.6	15
2	2021-11-02 12:55:16	378.9	30.8	15

Illustration 9: Example of the material tracking data Ⓓ. Every observation consists of the ID of the tracker, the time, the measurement is taken, and the coordinates in x- , y- and z-direction of the location of the tracker.

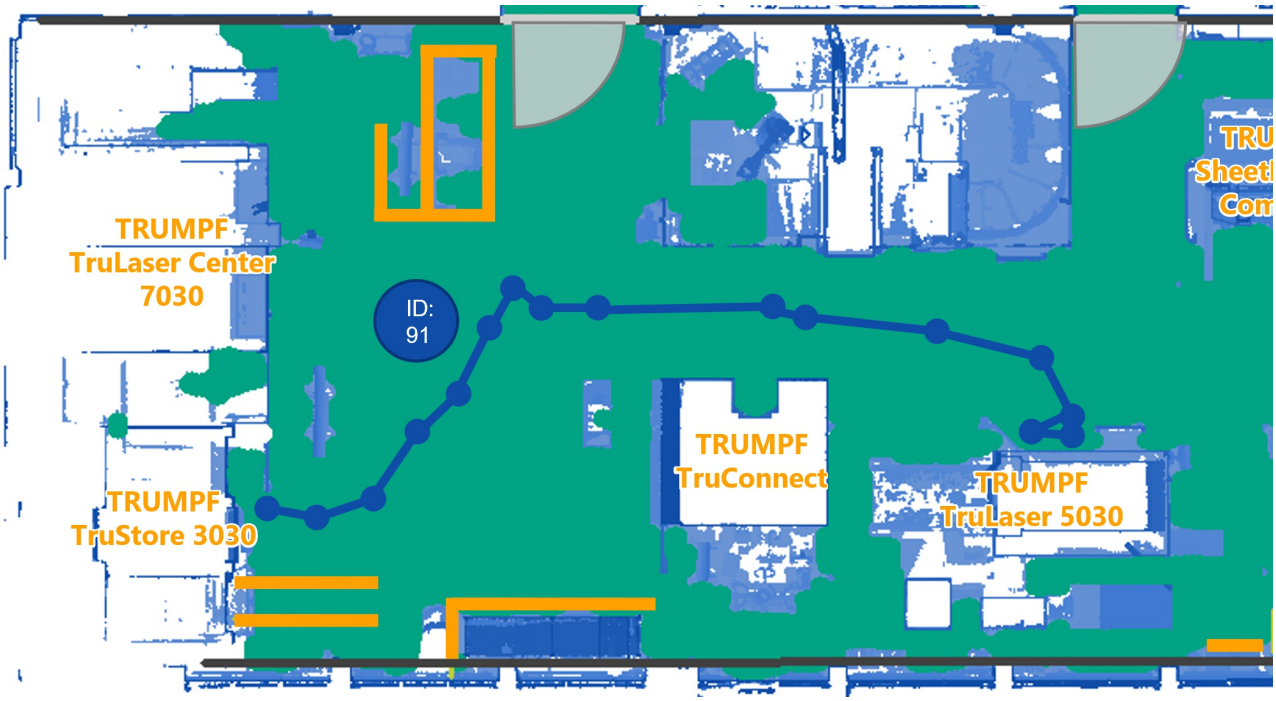

Illustration 10: A trace shows the movement of a material in the shop floor. In this example, material was taken out of the storage and transported to the laser machine. The ID of the tracker is 91. With the movement of the material, also the movement of a single employee responsible for the transportation becomes apparent.

and the coordinates constitute the minimal set of information required to solve the task.

Risk

- For the shop floor owner the existence of the material tracking data Ⓓ bears the risk that this data and the inherent business sensitive information will be used in an undesirable way.
- As no information about the service provider is contained in the material tracking data Ⓓ the service provider is not exposed to risks by the data.

Use case inherent risk mitigation

In this use case, the processing of the material tracking data is automatically done without a human in the analytical loop.

[Who] This reduces the involved parties, that is the number of potential criminal actors.

[How] An automated process limits the opportunities to use the data in an undesirable way.

Moreover, the data are not visualized, and from the raw data humans can only extract simple information such as the number of trackers in use.

Aspects of negotiation

[Data usage] It needs to be agreed upon whether the service provider is allowed to use the material tracking data Ⓓ only for the simulation or whether other usages are possible as well. The shop floor owner can use the material tracking data Ⓓ for every use case wanted. In this case an adequate compensation for the costs of generation might be paid by the shop floor owner to the service provider.

[Storage] Should the material tracking data Ⓓ be stored? At which point in the data pipeline? Where?

[Access] If the data are stored, who – person or process – should have access?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the service provider is only allowed to use the material tracking data Ⓓ to infer the traces of the material in the shop floor.

- the data are not stored at all and deleted directly after processing.
- only the process to infer the traces has access to the data.

2.2.2. ANALYSIS OF PERSONAL INFORMATION – MATERIAL TRACKING DATA

In subsection 2.2.1 the risk-value assessment of data pipeline [2] for business sensitive information was described. This subsection performs the risk-value assessment for personal information.

As already explained for data pipeline [1], personal information is not required nor wanted to run the simulation for the shop floor's optimization. The simulation requires purely factual inputs about layout, machines and processes. It is therefore neither necessary nor useful to generate personal information and to provide it to the simulation.

Personal information

[Direct] Personal information are not present directly in the material tracking data Ⓓ.

[Indirect] But personal information can be inferred indirectly. For example, the tracking data timestamps can be cross-referenced with a staffing plan of the employees. Doing so, a single trace could be linked to a specific employee moving the trolley. A visual example of a trolley's path is presented in illustration 10.

Value

This indirect personal information is of no value for any stakeholder. It is not useful for the purpose of optimizing production processes.

Risk

There is a risk that this indirect personal information could be misused. Particularly, employees are at risk of being benchmarked by the shop floor owner. For example, by pre-processing the tracking timestamps and coordinates, the speed of the workers is inferable.

Whether such benchmarking is generally permitted or prohibited depends on the respective employment contracts or existing company agreements. In any case, the employee should be informed by the shop floor owner about all performance evaluations that have been carried out.

Use case inherent risk mitigation

The risk of misuse of this personal information, however, is low. The processing of the tracking data in this use case is automatized, with only minimal manual intervention for trace validation. The paths are later used to infer production processes. The single traces can be deleted after the processes are identified. Depending on the agreement about access to the traces between the service provider and the shop floor owner, the latter might never have access to the traces at all. The probability of misuse, given the limited opportunities to extract the traces and link them with other sources, is rather low.

Aspects of negotiation

There are a few aspects of negotiation to further minimize the risk of personal information being potentially misused:

[Access] Does the shop floor owner have access to the traces?

[Data minimization: timestamps] Is it possible to omit the timestamps or "blur" the time information in such a way that the traces are still recognizable, but no connection can be established between the trace and the employee co-generating it?

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the shop floor owner has no access to the single traces.
- the single traces are deleted after the processes are identified.
- the timestamps are blurred appropriately.

The results of the negotiations should be included in the service contract signed by both parties.

2.2.3 SUMMARY DATA PIPELINE [2]: MATERIAL FLOW

The material tracking data Ⓓ contains business sensitive information about the shop floor owner. If one wants to carry out the simulation, this sensitive information is indispensable. The risk for the shop floor owner can be mitigated by a number of measures, but not eliminated.

The material tracking data Ⓓ contains indirectly derivable personal information about the employees of the shop floor owner. Since personal information are not

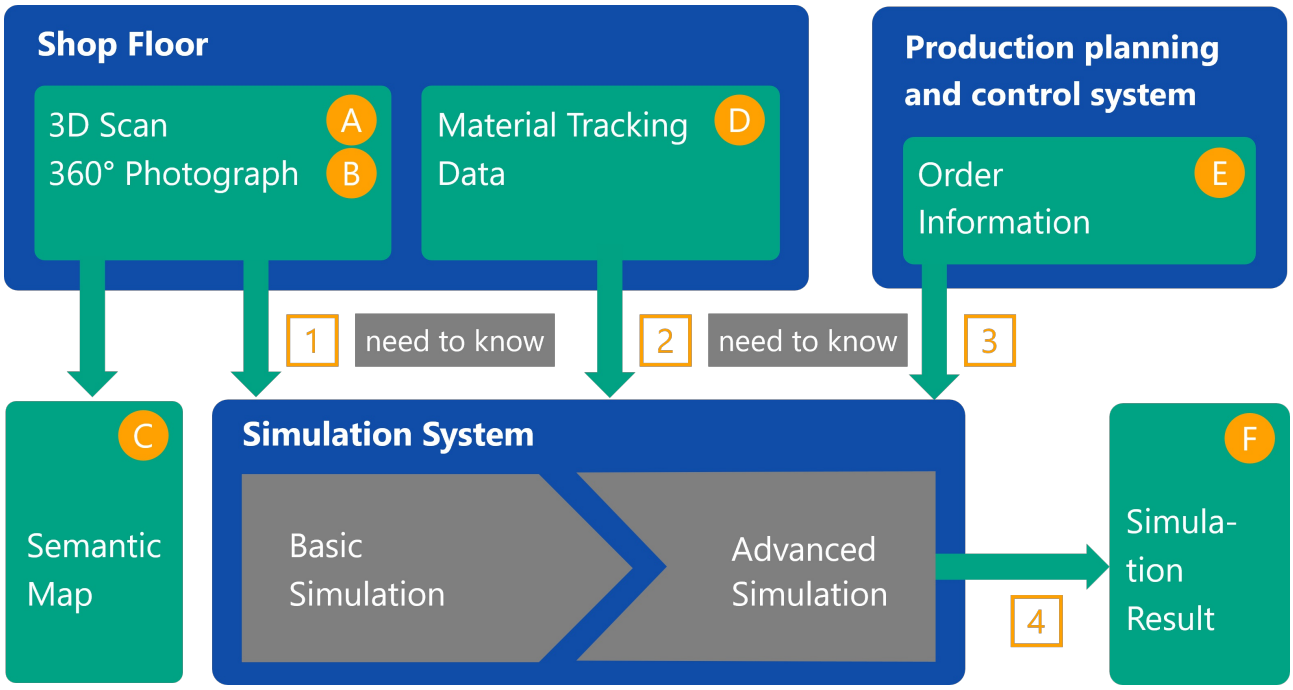

Illustration 11: Update of the data pipeline illustration 8 incorporating the results of the analysis of data pipeline 2: Only information the service provider needs to know to run the simulation should be provided.

relevant for the simulation, it is advisable to take all available risk mitigating measures with regard to personal data.
Analogous to the recommendations for data pipeline 1, it is generally recommended to strictly implement the need-to-know principle. This is shown schematically in illustration 11. Neither should more data be collected than is required for the implementation of the use case (data minimization), nor should access to the data be granted to anyone or any process not necessary to run the simulation. It is also advisable to delete the material tracking data D once the analysis has been completed.

2.3. DATA PIPELINE 3: ORDER INFORMATION

The order information pipeline 3. processes order data extracted from a production planning and control system. In this production planning and control system, the details of all orders are stored along with lots of other information. These order details include, for example, the name of the customer, the time of the order, any express options, the selected material and certainly the scope of the order. Especially the material properties are key to effectively understand – and afterwards optimize – the production process. The data source E contains a subset of this massive database. The information in E is selected from the system depending on its added value for the planned process optimization.

The production planning and control system and its database is hosted and maintained by the shop floor owner. As a subset of the overall database, data source E is also under the control of the shop floor owner. Hence, the order data E will be provided to the service provider by the owner.

The analysis of business sensitive information and personal information in the data will be carried out next.

2.3.1. ANALYSIS OF BUSINESS SENSITIVE INFORMATION – ORDER DATA

For the simulation, the order data (E) need to be extracted from the database of the production planning and control system. This extraction can be critical in itself and involves risks. This aspect will not be discussed further. Rather, in the analysis of the present use case, it is assumed that the extracted data source (E) is transferred digitally to the service provider via a technical interface.

Whether – and to what degree – business sensitive information is contained in the data source (E) depends on which information is specifically extracted. For the simulation, it is irrelevant who the specific client is. On the other hand, the material of a workpiece, and thus its weight, is of great importance for the simulation.

Information can be extracted directly from the order data (E) either by a human being or by a technical process. The information contained is easily accessible due to their descriptive nature. With this easy-to-acquire understanding, conclusions can be drawn, for example that 37.41% of all orders have an express option or that an average of 13 workpieces are ordered per order. It is undisputed that this is business-sensitive information.

Through further processing, in particular the fusion of the order data (E) with the material traces (D), further business sensitive information can be extracted. For example, it can be seen from the merged data that the production process for thin sheets up to a thickness of 2.3cm regularly comes to a standstill due to an under capacity for bending. Exactly this business-sensitive information is at the core of the insights to be gained through the simulation in this use case.

Business sensitive information

The provided order data (E) is business sensitive.

[Direct] The data contain direct business information.

[Indirect] The data contain indirect business sensitive information revealed for example after fusion with the material tracking data (D).

Value

The order and material information are of high value for the production optimization.

[Essential] The order data are essential for process optimization. For example, knowing the weight of a work piece is essential to correctly understand the needs of the production process.

[Data minimization] The specific metrics contained in the order data can differ significantly depending on the production planning and control system. Differently from other analyzed data sources, the actual data structure of the order data is unknown. Thus, whether the available data form a minimal set of information regarding the simulation is unknown.

Risk

- For the shop floor owner the order data (E) pose a serious risk.
- The order data (E) are free of risk for the service provider as they contain no information about the service provider.

Use case inherent risk mitigation

Since in this implementation the shop floor owner himself exports the data from his system, he has full control over the volume and variety of the information passed on. If he does not want to share a data point, the service provider has no possibility to enforce the release or to generate the information himself.

Aspects of negotiation

[Data usage] The shop floor owner has full control over the order data as it is all contained in the production planning and control system. He is thus able to use them in any way. It must be negotiated whether the service provider may use the data for purposes other than simulating the shop floor.

[Data minimization] It must be determined which information is to be specifically extracted from the production planning and control system.

[Storage] Should the data source (E) be stored separately

[Access] How exactly should the service provider access the data?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the owner should only provide strictly needed information.
- the order data (E) are not stored after being fed into the simulation.
- no access other than for the simulation is granted. Especially no access is granted to third parties.
- the data are only used for the simulation of the shop floor.

2.3.2. ANALYSIS OF PERSONAL INFORMATION – ORDER DATA

In subsection 2.3.1 the risk-value assessment of data pipeline [3] for business sensitive information was described. The present subsection performs the risk -value assessment for personal information.

As already mentioned, personal information is not required nor wanted to run the simulation for shop floor optimization. The simulation requires purely factual inputs about layout, machines and processes. It is therefore neither necessary nor useful to generate personal information and to provide it to the simulation.

Personal information

[Direct] If the order data (E) contain information about the client, personal data is contained directly.

[Indirect] It certainly depends on the exact subset of the data extracted from the system whether any form of indirect personal information is included. For example, it could be possible to derive which employee entered the respective data, or which employee is responsible for a specific customer. This aspect should be looked at again when the specific data set is available. At the moment, only conjectures can be made. These are omitted here.

Value

The client information potentially extracted from the production planning and control system is of no value for the purpose of the simulation.

Risk

If customer information is contained in the order data (E), there is a risk that it will be used in an undesirable manner.

Use case inherent risk mitigation

Since in this implementation the shop floor owner exports the data from the system, client data can be easily excluded from the export.

Aspects of negotiation

It should be noted in the service contract that the shop floor owner does not pass any sensitive information about his customers on to the service provider.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- no information about the clients of the shop floor owner is contained in the order data (E).

2.3.2. SUMMARY

The order data (E) contain business sensitive information about the shop floor owner to the highest degree. Since the order data play a central role in the simulation they cannot be omitted. If the shop floor owner is interested in the simulation, the risk has to be taken. Measures to mitigate the risk tend to be limited to measures for storing and accessing the information.

Whether or not the order data (E) contain personal information depends on the information extracted from the production planning and control system. Personal information can be omitted entirely by not exporting it at all. Since this personal information does not add any value to the simulation this can be done without a loss of quality for the simulation.

Here, too, it is advisable to implement a strict need to know principle: only the data relevant for the simulation should be extracted and passed on to the service provider. After the simulation is completed, the data should be erased. At the same time, access to the data should be limited to what is absolutely necessary. The recommendation for data pipeline [3] is shown in illustration 12.

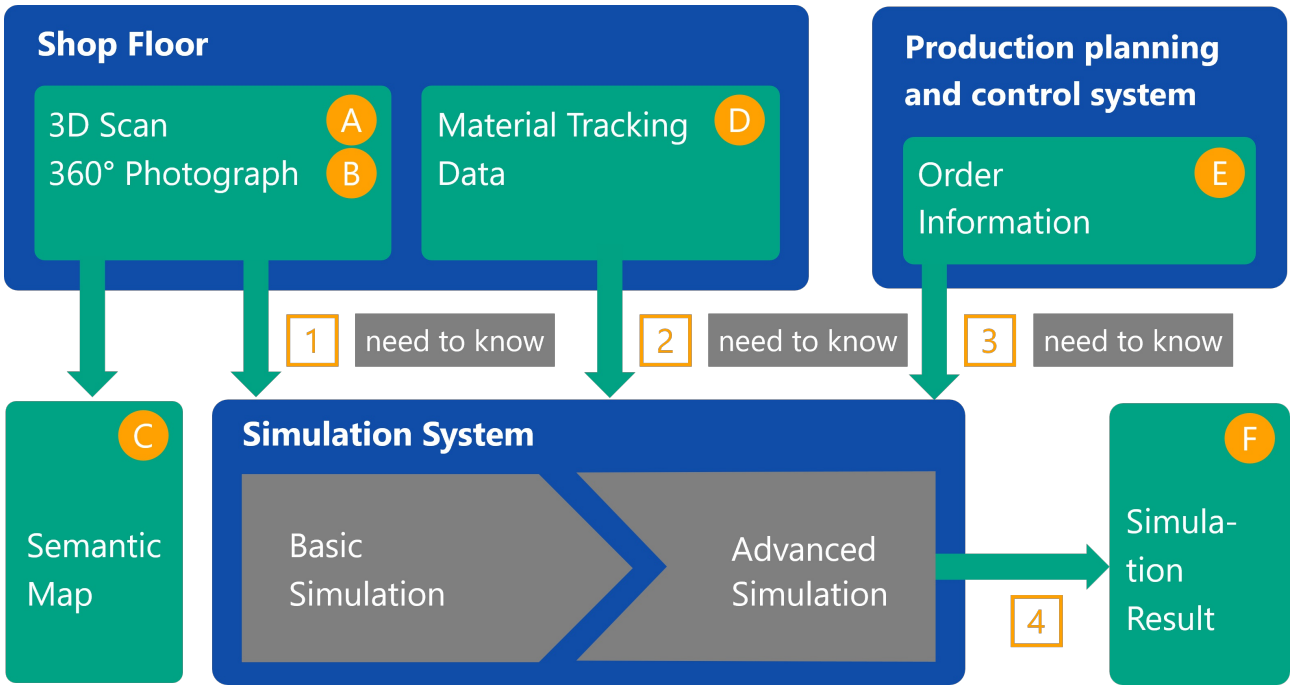

Illustration 12: Like for the previous data pipelines, for data pipeline 3 only information strictly required to run the simulation should be delivered to the service provider.

2.4. DATA PIPELINE 4: SIMULATION

In this use case, the data pipeline 4 processes the outputs of data pipelines 1, 2, and 3 automatically. It performs a simulation using data sources C, D, and E. The output of the data pipeline 4 is the simulation result F. The result can be numerical data as well as visualizations. It is used to recommend changes in the shop floor to optimize production processes.

There might be sensitive information involved. The presence of business sensitive information and personal information is discussed next.

2.4.1. ANALYSIS OF BUSINESS SENSITIVE INFORMATION – SIMULATION RESULTS

The shop floor owner, more precisely his entire production facility including all production processes and procedures, are completely transparent through the results of the simulation and optimization F. In particular, weak points in the production and potentials for improvement are identified. Nobody wants this information in the hands of a competitor or even a customer. The simulation results are sensitive to the highest degree.

How understandable and accessible the results F are depends on whether they are visualized or, in a more technical form, pure data points. It can be assumed that the service provider prepares the results accurately and in a comprehensible manner.

A precise description of the shape and form of the simulation result F is currently not possible. There will most likely be a visualization of the shop floor comparable to the semantic map C as shown in illustration 5, in which processes and also conspicuous observations are shown. The potential for improvement will probably also be made available in one or more graphical representations and stimulate discussion.

How the simulation result F is made accessible is a crucial driver for the risk profile – both for the shop floor owner and for the service provider. This is illustrated by two aspects:

1. Are the results F only available within the service provider's simulation environment or also outside? If the results are only available in the service provider's simulation environment, the service provider can control the scope and thus the op-

portunities for insight. If, on the other hand, the results are available to the shop floor owner outside of the simulation environment, he can deal with the results without being influenced by the service provider. Furthermore, the shop floor owner has the opportunity to include third parties in the evaluation without involving the service provider. Without the guidance of the service provider the risk of misinterpreting the results occur.

2. Are the results only accessible in a visually processed form or is the underlying data also made available? If the results are available in a visualized form, the possibilities of gaining knowledge are constrained by the visualized elements. If the results themselves are available as raw data in digital form, there is the possibility of "result" mining. Not only can the situation on the shop floor be inferred from digital results, but it could also be possible to infer properties of the simulation used. There are business sensitive information from the perspective of the service provider.

Business sensitive information

[Direct] The simulation results Ⓕ contain direct business information of the shop floor owner.

[Indirect] The simulation results Ⓕ certainly contain also indirect business sensitive information. This includes:

- It can be assumed that further business sensitive information of the shop floor owner can be inferred indirectly from the simulation results Ⓕ. However, further processing or linking of the simulation results Ⓕ is not subject to the use case.
- It might be possible to derive details about the process optimization procedure by applying reverse engineering on the simulation results Ⓕ and thus, revealing business sensitive information of the service provider.

Value

[Essential] The value of the entire use case stems from the great importance of the simulation results Ⓕ. This enables the shop floor owner to optimize his production. If the results Ⓕ were not essential, the entire use case would be obsolete.

[Data minimization] It should only be shown what really affects the production facility and the production processes of the shop floor owner. The extent to which not required information is available depends on the negotiated input data sources Ⓐ, Ⓑ, Ⓒ, Ⓓ, and Ⓔ. If, for example, the condition of a machine has not been inferred, it should not appear unnecessarily in the simulation results Ⓕ.

Risk

- The generation of the simulation results Ⓕ reveals large parts of the production processes and the weaknesses of the shop floor. The unwanted use of this business sensitive information can significantly harm the shop floor owner.
- By processing the simulation results Ⓕ in an unwanted way, the service provider's process optimization methodology can be reverse engineered. A competitor could use this information to build or improve a similar service.

In contrast to the data sources Ⓐ, Ⓑ, Ⓒ, Ⓓ, and Ⓔ, here both stakeholders have a significant risk for revealing business sensitive information.

Aspects of Negotiation

[Data usage] It needs to be agreed whether the service provider is allowed to use the simulation results Ⓕ for other usages. On the other side, it needs to be negotiated, whether or not the shop floor owner is allowed to share the results with third parties.

[Connection to simulation tool] It must be clarified whether the results should only be available within the service provider's simulation environment or also outside of it.

[Type of representation] It must be clarified in which form the results are made accessible: visualized and/or as raw data.

[Storage] Which parts of the simulation results should be stored? Where?

[Access] Who should have access to what part of the simulation results Ⓕ?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

The risk inherent to the use case can be minimized if

- the data is not shared with any third parties.

- it is not allowed to use the results to any other course than to derive insights about the shop floor processes.
- there are clear rules, if and where the results are stored.

2.4.2. ANALYSIS OF PERSONAL INFORMATION – SIMULATION RESULTS

The simulation results (F) should not contain any personal information for two reasons:

1. Nothing of interest in the context of optimizing the processes and the layout of the shop floor constitutes personal information.
2. The level of abstraction, for example from the single traces of the material flow to processes, eliminates personal information.

Personal information

[Direct] There is no direct personal information.

[Indirect] The simulation itself and the output should generally not contain indirect personal information. However, the material tracking might be part of the simulation output. In this case personal information is indirectly contained (refer subsection 2.2.2.).

Whether in visualized form or as technical data, there is one reason why the material tracking data (D) could be part of the simulation results: phenomena in the production processes might need to be validated. For the validation, single traces displaying the actual moving on the shop floor might be required. Certain challenges for the production flow might be best illustrated with one trace that shows the pattern of interest.

Value

As a part of the simulation results (F) the material tracking data (D) are valuable. They support validation and explanation of concrete findings within the process flow.

Risk

There is a risk that this indirect personal information is being misused. The risk replicates the risk described for material tracking data in subsection 2.3.2.

Use case inherent risk mitigation

The material tracking data (D) play an important role in explaining and understanding the results of the simulation. However, this does not imply that all material tracking data (D) must become part of the simulation results (F). Rather, the service provider can select a few prototypical material traces for illustration. The date and time information of these illustrative traces can be omitted.

Aspects of negotiation

[Data minimization] What is really necessary for the shop floor owner to utilize the simulation results?

[Storage] Should the material traces be stored as a part of the result (F)?

[Access] Who should have access to the stored material traces?

The results of the negotiations should be included in the service contract signed by both parties.

Recommendation to minimize overall risk

Assuming that the material tracking data is indispensable for understanding the results, only a few traces required for the explanation should be made available by the service provider. The date and time at which these traces occurred should not be specified. The exemplary traces can be deleted as soon as they were able to understand a conspicuous situation.

2.4.3 SUMMARY

As result of the data pipeline [4], the results (F) contain both business sensitive information about the shop floor owner and, in indirect form, about the service provider.

- The results (F) show the strengths and weaknesses of the layout of the shop floor and the production processes as well as the potential for improvement. From the point of view of the shop floor owner, it is essential to control and minimize the group of people who know these results. Passing it on to unauthorized persons can result in significant damage for the shop floor owner. For risk mitigation the shop floor owner can only control the storage of and the access to the results (F).
- Risk arises for the service provider if the results (F) are transferred in digital form to the shop floor owner. In this case it may be possible to reverse engineer the methodologies used in the simulation. The service provider can significantly minimize its risk by not making the results available in

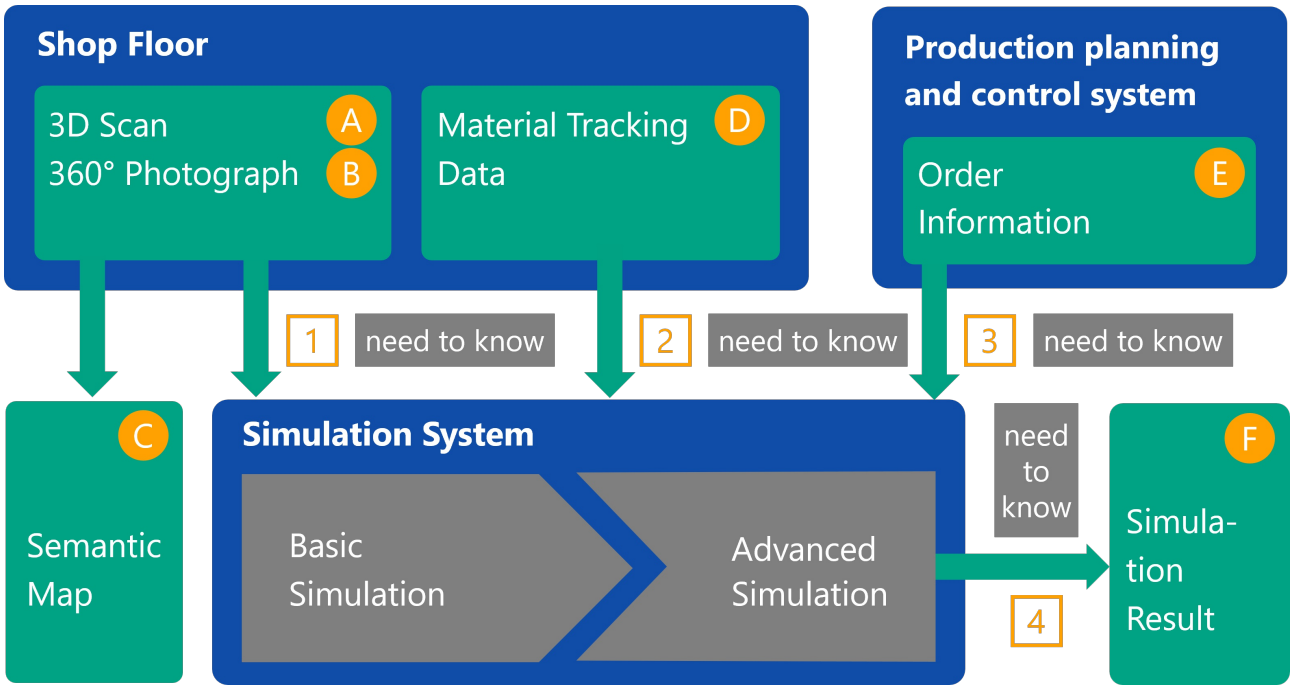

Illustration 13: Overview about the final system of data pipelines relevant to generate the simulation results F. Given the presence of business sensitive and personal information in the process and the results, the data should be handled with care. Risk mitigating measures should be taken wherever possible. Only those data should be generated, processed and provided that are necessary to run the simulation. Valuable other results, like the semantic map C, should be seen as separate products with own terms and conditions.

digital form and not outside the simulation environment.

Personal information is only part of the simulation results F if the material tracking data D is passed on as part of the result. The material traces can play a role in explaining and understanding the findings of the optimization. The risks for the employees of the shop floor owner can be mitigated by minimizing the number of exemplary traces provided and by removing the date and time information to make it more difficult to infer the employee.

As shown in illustration 13 the need-to-know principle should be implemented for the data pipeline 4: The simulation results F should not contain irrelevant information and the access to F should be limited to the minimum.

3. CONCLUSION AND OUTLOOK

Thank you for reading this far.

3.1. CONCLUSION

In this paper, an approach to negotiate adequate data processing, based on a risk-value assessment for sensitive information, is presented. Sensitive information encompasses both personal and business sensitive information, where the latter could be intellectual property, trademarks, or business secrets. While regulations are more concerned with the protection of personal information, for companies however, business sensitive information are just as important to safeguard. In any case, it is essential to manage this sensitive information in a fair and transparent way. Fairness does not necessarily translate to equal distribution: Not everything has to be accessible to every stakeholder. A fair data access should aim at maximizing potential value while minimizing risks.

The advantage of cyber-physical systems is that the co-generation normally materializes after a contract is signed. Thus, negotiations take place in any case. The only enhancement needed is to negotiate the aspects around data as well.

To carry out the risk-value assessment, the simulation service is first broken down into the independent data pipelines it consists of. Every data pipeline, more precisely every data source processed within the data pipeline, is analyzed for the content of business sensitive and personal data. If sensitive information is present, it has to be clarified whether this information content is unavoidable. If this is the case, risk mitigating measures are presented. The most important element is that the two stakeholders negotiate how the data is to be handled in detail. Due to the co-generation, both stakeholders have a say.

In the use case presented, personal data plays no role in the simulation. They are neither needed nor wanted and can be greatly reduced with a little care. In particular, taking the scans and photographs in an empty shop floor helps greatly.

Since the aim of the simulation – the optimization of the process flows – is highly sensitive for the shop floor owner, it is not surprising that all of the data sources contain business sensitive information. For the protection of the shop floor owner, it is important that the need-to-know principle is strictly adhered to and, ideally, no data is stored. In any case, no third party should have access to the data, and the service provider should only be allowed to use the data to perform the simulation.

Of course, the contractual partners are free to agree that the service provider may use all, or one of the intermediary data sources, for other purposes. For example, to improve its product. What the shop floor owner receives from the service provider in return is up to the two parties.

3.2. OUTLOOK

The risk-value assessment allows in a structured and unemotional fashion to develop a cyber-physical system where the handling of data is done in a way that maximizes the value added by controlling the arising risks. This approach is not restricted to cyber-physical systems but can be applied to any system, especially if more than one stakeholder is involved.

Effort that pays off

The discussion in this paper may give the impression that a risk-value assessment involves a great deal of effort and is therefore a great burden.

This may be the case when such an assessment is being conducted for the first time by one or more stakeholders.

However, it can be assumed that clear and simple standards for the procedure will emerge quickly. The recurring structure of the risk-value assessment shows that the same aspects are always important. What sensitive information content is present? Can it be reduced without making the desired evaluation impossible? Is the data stored? Who has access?

While this procedure for personal information is now routine, there is still a lot of catching up to do when it comes to business sensitive information.

Change of status quo required

Basically, a change in the mindset of all stakeholders is required:

- Do not blindly buy a service that uses your data without considering the potential impact of data processing.
- A service that generates and processes sensitive data is not offered as a black box, but instead presents the value and the risks transparently.

One effect of the risk-value assessment and the fact that the handling of the data requires in many cases an adjustment to the status quo is that systems have to be adjusted in their structure. Currently there may be one or the other machine that loads all the co-generated data to a cloud. However, it can be assumed that not all the data, but only a subset, belongs in this cloud after negotiations have taken place. If the cloud is operated by the service provider, risk-mitigating measures must be taken that minimize the information content regarding business sensitive information of the user as much as possible. One consequence is that, for example, data must be aggregated on a machine and only these aggregates are allowed to be loaded into the cloud. Furthermore, the shop floor owner could (rightly) request to receive the data generated by the machine. This would require another interface on the machine. With this data stream, care should again be taken to ensure that as little as possible of the service provider's IP goes to the user.

Market advantage

Finding good solutions for handling data and processing it in the CPS environment, setting up hardware, software, and data processes sensitively, together with openness to negotiations will bring providers a significant advantage on the market in the medium term at the latest. The more burnt the customer already is, the lower his appetite for risk, the more advantageous is a product that enables transparent data use in the B2B environment with and without CPS.

Concessions to practice required

In practice, a "totalitarianism" that preaches: "Thou shalt never save a date" does not make sense.

In the use case shown, it is not so much the hardware that needs to be adapted, rather the processes for handling the data. Dissecting the overall process into multiple data pipelines helps to understand how precisely to distinguish between the data to be processed, the data to be passed on, and the data to be stored. In practice, you will probably only delete a 360° photograph or a 3D scan when it can be seen from subsequent process steps that the processing of the 360° photograph or 3D scan was successful. Nobody wants to repeat a complex scan of a shop floor because a program had a bug or the calculation had to be aborted due to a hardware problem. But that does not mean that you can simply throw all the data into one big pot that you can access at any time and on any occasion.

Don't be afraid to deal with sensitive information

Many companies currently believe that they are not allowed to do anything with personal information. The main focus of this use case was not on this aspect. However, the discussion above shows that it is in no way illegal to work with personal data. It is important to have consent and to clearly define what is processed and for what purpose. The situation with business sensitive data is completely analogous. Not everything is forbidden. But not everything is allowed.

The positive aspect of a B2B situation is that there is always a contractual agreement. The contractual partners have the opportunity to negotiate at eye level how the co-generated data should be processed and who should have access to which information. Once this has been done, the consent and the possibilities for data usage are clearly given. These agreements should become part of the service contract, which increases the legal certainty for both parties.

Challenge indirect sensitive information

With every risk-value assessment and with the negotiations about how to handle the data, the big challenge will always be to identify the indirect information content. What potential does the data have? Such a question can never be answered definitively. It remains to be seen what an archaeologist of the future will be able to learn about our age from these data. What insights could a terrorist glean from the data? It is just as impossible to identify all conceivable data usage options as it is to determine which sensitive information could come to light through these usages. Therefore, a residual risk always remains for all parties involved in data generation.

In the B2B context, the key stakeholders who have to bear risks and who could be exposed to potential damage are known. The direct and indirect possibilities of gaining knowledge from the data in the respective business environment can certainly be worked out. These possibilities will also be very similar for different stakeholders: the layout of the shop floor is business sensitive information for every shop floor owner. The risk that a competitor might learn something from the layout is there for everyone. The only difference is whether one sees a smaller or larger problem in this circumstance.

A list of standard situations in which data processing reveals sensitive information will therefore emerge quickly. These standard situations are likely to represent the vast majority of the basic and realistic usage options, ignoring science fiction scenarios and potential criminal acts. The only task left between the respective stakeholders is to clarify on this basis, whether risk mitigation measures should be taken, and if so, which exactly.
As soon as this has happened, a regular review process is sufficient to clarify whether new technologies, new processes, or new machines have changed the situation and thus the risk profile. If there is a change in the risk profile, new agreements can be made and/or old ones revised.

Increase awareness of opportunities

It is in the nature of a risk-value assessment to address primarily the risks. The list of standard situations just mentioned can also be interpreted the other way around: as a list of chances! Not everyone involved is aware of the added value that can actually be drawn from the available data. The discussion within the framework of the risk-value assessment encourages those involved to think about what else might be possible for them.

Strengthening European values

At first glance, the risk-value assessment for sensitive data looks like a millstone around the neck of those involved:

- The service provider will be restricted. He can no longer simply manage and use all data in his sovereignty, as is often the case in current implementations.
- The shop floor owner has to deal much more specifically with the ordered product or service. Agreements have to be made. Contract clauses have to be defined.

But the other side of the coin is that trust can evolve. Should this trust be broken, both stakeholders have legal recourse thanks to the very clear contractual provisions. This legal certainty also supports the development of trust.
On this basis, the digital gold – the data – can be used with added value for all sides, without one side cheating the other, without the rights of employees being ignored, and without European values being called into question.
Rather the presented approach supports the defense of European values in the digital realm.

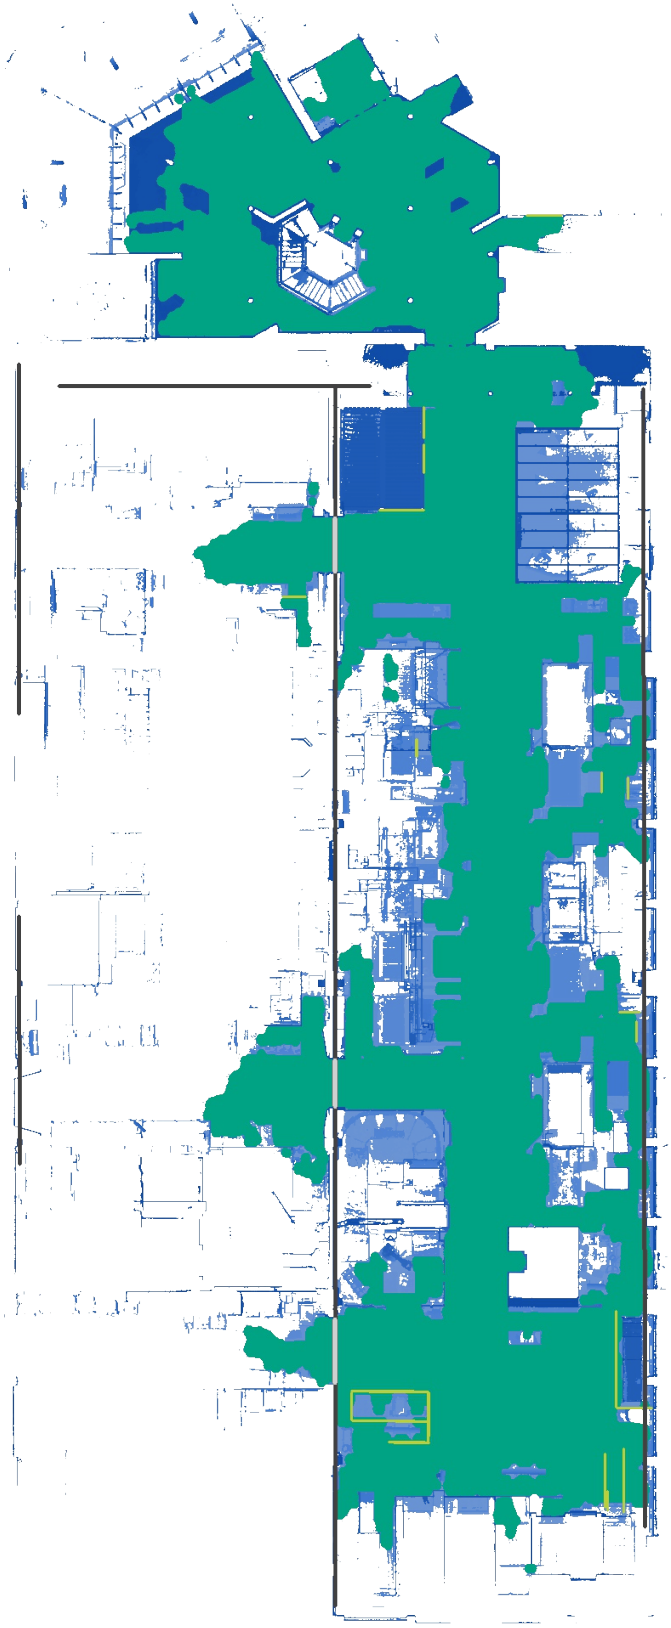

Illustration 14: Example of shop floor visualization.

ISSN: 2941-1769

acs plus GmbH
Rahel-Hirsch-Str. 10
10557 Berlin
www.acs-plus.de

acs plus
data with care