Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

HOW NOT TO VISUALIZE YOUR SIMULATION OUTPUT DATA

Jonas Genath
Steffen Strassburger

Group for Information Technology in Production and Logistics
Ilmenau University of Technology
Max-Planck-Ring 12
Ilmenau, 98693, GERMANY

ABSTRACT

Hybrid modeling and simulation studies combine well-defined methods from other disciplines with a simulation technique. Especially in the area of output data analysis of simulation studies, there is great potential for hybrid approaches that incorporate methods from machine learning and AI. For their successful application, the analytical capabilities of machine learning and AI must be combined with the interpretive capabilities of humans. In most cases, this connection is achieved through visualizations. As methods become more complicated, the demands on visualizations are increasing. In this paper, we conduct a data farming study and delve into the analysis of the output data. In doing so, we uncover typical errors in visualizations making the interpretation and evaluation of the data difficult or misleading. We then apply concepts of visual analytics to these visualizations and derive general guidelines to help simulation users to analyze their simulation studies and present results unambiguously and clearly.

1 INTRODUCTION AND MOTIVATION

Hybrid modeling and simulation studies follow the idea to use simulation as a connecting link between methods from different research disciplines (Mustafee and Powell 2018). Especially in the area of output data analysis, there is great potential to use hybrid approaches that incorporate data analysis, machine learning and artificial intelligence (AI) methods (Feldkamp et al. 2020) to mutually create greater value for the analyst (Tolk et al. 2021).

One approach to go beyond traditional simulation studies, which usually aim to perform scenario-based analysis or even simulation-based optimization, is data farming developed by Horne and Meyer (2005). Data farming aims to understand the behavior of the model in terms of the relationship between factors and outputs, potentially uncovering knowledge and interesting relationships that were previously hidden (Feldkamp et al. 2015). Data farming refers to the method of extracting large amounts of output data from the simulation model by using large-scale experimental designs, high-performance computers for massively parallelized experiments to focus on more complete coverage of possible system responses and machine-assisted analysis (Horne and Schwierz 2008). The research within the field of data farming has also always been concerned with the application of advanced data analysis methods to process these large volumes of simulation output data and produce appropriate insights (Lucas et al. 2015; Sanchez 2014). The concept of Knowledge Discovery in Simulation Data (KDS) was developed to delve deep into the analytics side of data farming and provide a detailed process model for applying white-box machine learning and data mining methods to large-scale simulation data (Feldkamp et al. 2020). But more methods are constantly emerging in the field of machine learning and AI, which can be combined with simulation though the concepts of data farming and KDS to derive new insights about the simulation model. These methods are often also becoming more complicated. For example in recent papers we extended the KDS concept with

979-8-3503-6966-3/23/$31.00 ©2023 IEEE

methods from explainable artificial intelligence (XAI) to enable the use of black-box algorithms like artificial neural networks and we showed that we could generate new and previously hidden knowledge about the system through these methods (Feldkamp 2021; Feldkamp et al. 2022). However, we also learned that the visualizations of the results by the XAI methods are not always intuitive or easy to interpret.

To derive the most knowledge from your data farming output data, the strengths of humans and utilized algorithms must be combined. This means that the analytical capabilities of machine learning algorithms must be connected with the interpretative capabilities and analytical reasoning of the simulation user. The interface between these two aspects are visualizations, but their complexity often also grows with the amount of data that is available and the complexity of the analytical methods (Keim et al. 2008). The research discipline focusing on this interface is visual analytics. However, most recent research in visual analytics focuses on developing custom software to solve specific problems (Cui 2019; Ham 2010; Matković et al. 2018). But especially in research, to develop or modify a visualization software every time for a specific problem is usually not efficient, because individual visualizations are needed, which allow the analyst or researcher the specific freedom to present their results appropriately. This is why many resort to programming languages to create their visualizations and share their research results. As an author or domain expert it is usually easier to interpret and understand the visualizations although for others this is not always the case. As a result, time and scientific or industrial opportunities can be lost due to information overload or inconclusive presentation of the results because a reader may not understand them correctly (Keim et al. 2008). In addition, Cui (2019) showed in his literature review on visual analytics that there are very few publications in the context of simulation. Furthermore, the simulation community is missing scientifically published guidelines and techniques for presenting results, particularly in relation to data farming, which can lead to inconsistencies and difficulties in interpreting and comparing simulation analysis results. However, the ability to communicate results accurately among peers is crucial for scientific discourse (Keim et al. 2008).

In this paper we want to make a first step in the direction of creating awareness for the importance of accurate and clear presentation of research results. To achieve this, we conduct an illustrative data farming study and delve into the analysis of the simulation output data using unsupervised learning. In doing so, we uncover typical errors in visualizations and apply concepts of visual analytics to them. Thereby we derive some method independent guidelines to help simulation users to analyze their simulation studies and present their results unambiguously and clearly. We also discuss general considerations and requirements for visualizations.

The remainder of this paper is structured as follows: In Section 2, we give an overview of the related work, namely being visual analytics and the combination of data farming and knowledge discovery in simulation data. Section 3 discusses the general considerations and requirements that apply to visualizations. In Section 4, we present the data farming study, by first introducing the scenario and simulation model, followed by the analysis and visualization of the simulation output data. We conclude the paper with some final remarks and a discussion of possible future work in Section 5.

2 RELATED WORK

2.1 Data Farming and Knowledge Discovery in Simulation Data

Simulation studies typically have a defined objective, and not having one is considered a major issue (Law 2003). In the past, simulation experts focused on minimizing computational effort due to limited computing time and memory space. Furthermore, in practical scenarios, simulation analysts may rely on their experience to make an informed estimate regarding which factors are likely to have a significant impact on achieving the objectives of the simulation study (Feldkamp 2021). However, Kleijnen et al. refer to this as a trial-and-error approach and argue that analysts should spend more time analyzing the model than building it (Kleijnen et al. 2005). Traditional, goal-based experimentation aims to answer specific questions, while the Data Farming method seeks to explore the entire range of possible system behavior to gain a better understanding of all potential options and system behavior (Horne and Schwierz 2008).

Data Farming combines large-scale simulation experiments with high-performance computing and big data analysis methods to maximize data output, similar to how a farmer cultivates their land to maximize his crop yield (Sanchez 2014; Sanchez and Sanchez 2017). Knowledge discovery in simulation data is an extension of data farming and provides a process model and workflow for using data mining and machine learning methods as well as interactive visualizations to analyze complex simulation models (Feldkamp et al. 2015). This is particularly beneficial for models that have numerous outputs displaying a complex, multidimensional response surface. These outputs may even be conflicting or in a trade-off situation with one another, making their mutual interpretation extremely challenging (Feldkamp 2021). By grouping outputs into categories and using multidimensional pattern recognition algorithms, the relation between factors and outputs can be investigated, as shown in Figure 1 (Feldkamp et al. 2020).

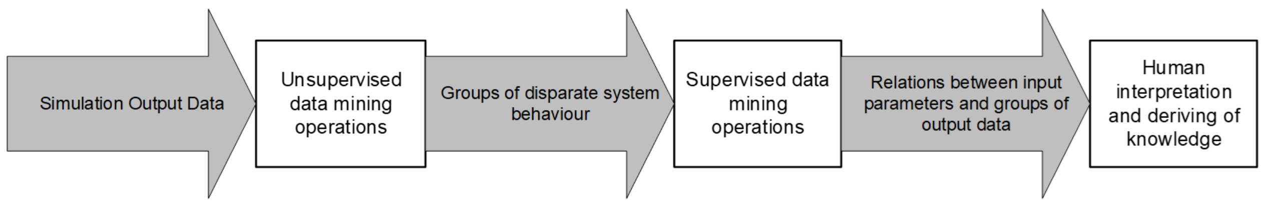

Figure 1: KDS process for analysis of simulation output data (Feldkamp et al. 2020).

The process shown makes it more manageable to analyze and interpret large amounts of simulation data. Various case studies have demonstrated the effectiveness of this approach, revealing previously unknown knowledge that would have been difficult to gain using traditional methods (Lechler et al. 2021; Strassburger et al. 2018). More precisely, discovering patterns in unstructured data is referred to as unsupervised data mining. Once the simulation data has been categorized into various groups of system behavior, the relationship between outputs and corresponding factors can be explored using supervised data mining algorithms. These algorithms create models that reveal the relations between simulation input and output data, which can subsequently be used to derive rules that contribute to knowledge creation through human interpretation. Essentially, every simulation experiment serves as a training record for a supervised algorithm, making it necessary to solve basically a classification problem (Feldkamp et al. 2020). While the KDS process provides recommendations on appropriate visualization formats for the corresponding data mining methods, it does not address the specific challenges or potential pitfalls that may arise when presenting the results visually. Visual analytics, which will be discussed in the following section, aims to overcome these challenges for presenting data mining results and complex data.

2.2 Visual Analytics

Visual analytics is a multidisciplinary research field that draws on visualization, advanced data analysis techniques, and analytical reasoning. Visual analytics leverages human perception and visual interactions to enhance the ability of researchers to gain valuable insights, discover new knowledge, and obtain a deeper understanding of large and complex datasets through effective data visualization (Cui 2019). Recent advancements in computing technologies and data visualization techniques have enabled visual analytics to become an essential tool for scientific research, business intelligence, and many other fields (Endert et al. 2017). Figure 2 illustrates how visual analytics incorporates a broad range of research fields and also depicts the integration of human cognition, perception, and reasoning with algorithmic data analysis in visual analytics and how visual analytics acts as an interface between humans and machines. Through interactive visual representations of the data, the perceptive skills, analytical reasoning, and domain knowledge of humans are coupled with existing data analysis processes to gain insight and new knowledge from data (Ellis et al. 2010). More precisely, visual analytics facilitates analytical reasoning by utilizing effective visualizations and interactive tools that enhance human capacity to perceive and explore complex data in a human-in-the-loop process, as discussed by Keim et al. (2010).

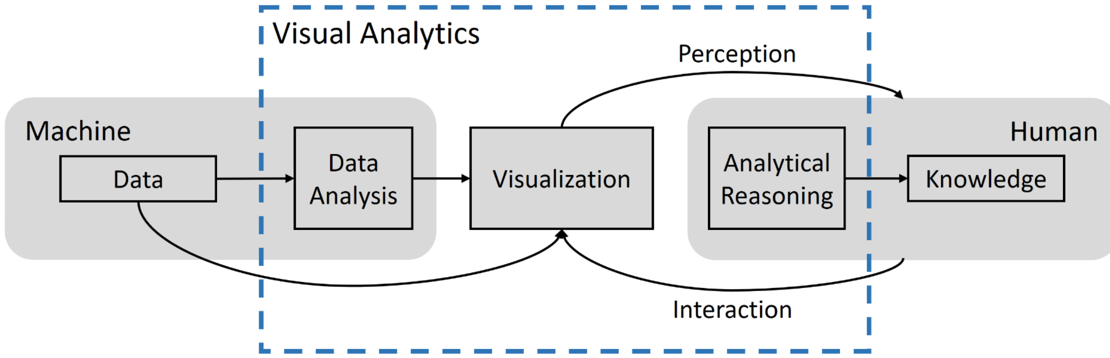

Figure 2: Visual analytics as the interface between humans and machines.

Specific visual analytics tools can be classified according to two different components: the dimension to be presented and the type of interaction (Cui 2019). Especially the interaction with visualizations is highly relevant for generating knowledge, but not necessarily for presenting that knowledge. Therefore, the primary focus of this paper is on the general visualization aspect of visual analytics and how it can enhance human perception, rather than on a specific technique or tool, which most of the time focus on the interactions. This paper aims to highlight common mistakes to avoid when visualizing data or developing a visualization software for simulation data. In the following chapter, we will outline the specific visual requirements that a scientific visualization should meet in accordance with the principles of visual analytics.

3 GENERAL CONSIDERATIONS AND REQUIREMENTS TO VISUALIZE YOUR DATA

Visualizations play a crucial role in understanding and interpreting simulation data. While visual aspects, such as color perception and imagination, may have a subjective influence, there are still general requirements and considerations that should be taken into account. Munzner (2014) and Ware (2021) provide the following considerations that should be considered when creating visualizations. First and foremost, visualizations must be accurate, properly labeled, and consistent with the underlying data. It is important to avoid introducing any bias or misleading representations that might obscure or misinterpret significant patterns or trends. This also means that the data should be presented in a way that accurately reflects its underlying meaning, with clear and appropriate labels for axes, scales, and units. Secondly, the choice of visualization technique should be appropriate for the data and the specific research question or hypothesis being investigated. This includes considerations such as the type (e.g., categorical or continuous), the scale (e.g., linear or logarithmic), the dimension, and the structure of the data (is the data sequential?). If these requirements are not met, then according to Keim (2001) and Spence (2014) a visualization is not expressive. Additionally, the visualizations should be designed with the audience in mind, considering the level of expertise and familiarity with the subject matter as well as the effort to create and interpret a visualization. This applies in particular to the analysis of data or the knowledge generation using complex or novel methods from the field of AI or machine learning. These considerations are needed to meet the requirements for an effective and appropriate visualization (Spence 2014; Ware 2021).

Although the point of interaction by the simulation user is crucial during the knowledge discovery phase according to the concept of visual analytics, it is difficult to implement this aspect in a scientific publication. Therefore, to preserve and facilitate the reader's comprehension of the iterative process, the procedure leading up to the presented results should always be described and highlighted through filtering or comparisons. For the reason of not overloading the viewer with unnecessary information, visualizations should also only be tailored to the specific goals and function of the analysis, as shown in Table 1. This means, the simulation user should always think about the goal of the visualization and how it can be achieved.

Table 1: Functions and goals for visualizations (derived from Spence (2014)).

Function	Goal
Identification	Highlighting various features of objects
Distinction	Delimitation of different objects.
Correlation	Revealing a direct relationship between objects
Association	Linking object relationships
Localization	Clarification of the relative or absolute distance between objects
Classification	Classify objects based on properties
Arrangement	Assigning a specific order of precedence
Comparison	Highlighting similarities or differences

Color schemes are an important aspect of visualizations and should be chosen carefully to ensure that viewers can easily distinguish between different categories or data points. It is important to avoid using colors that are difficult to distinguish. Colors can also be used effectively to achieve different goals in visualizations, as shown in Table 1. However, it is important to note that the more colors used in a visualization, the more complicated the analysis becomes for the viewer (Moreland 2009). Therefore, when something is to be highlighted, it may make sense to gray out all data points that are not relevant and not show each feature in a specific color. Creating and using color maps is also a good way to match colors and their effects. Brewer (2016) distinguishes between qualitative, single and multicolor sequential, rainbow, and divergent color maps. Sequential color maps run from a highly saturated color to a very unsaturated one, close to white. They are particularly well-suited for ordinal data and can be easily transferred to scalar values (Moreland 2009). For example, overlaps of data points can be represented by a more intense color. Diverging color maps consist of two main colors, which merge into each other by a third, but unsaturated color, such as yellow or white. Qualitative color maps are ideal for representing nominal data because their deliberately unordered sequence is good for distinguishing individual discrete data values or groups. The use of a rainbow color map is discouraged by Moreland (2009) and Schwabish (2021) because of disregard for perceptually specific requirements. Continuous values are difficult to map unambiguously because a uniform min-max value assignment is not possible, and the transitions of the colors are not perceived uniformly (Moreland 2009).

Finally, visualizations should be accompanied by explanatory context (Keim et al. 2008), such as a description of the simulation model and parameters, and an explanation of the specific analysis being presented. Overall, visualizations that effectively communicate their findings and insights to a wide audience require careful consideration of both the data itself and the intended audience, with a focus on accuracy, clarity, and appropriate design choices.

4 DATA FARMING STUDY

4.1 Scenario and Simulation Model

To analyze and illustrate the potential problems that can arise in visualizations, we conducted a data farming study using an academic use case to generate a sizable quantity of data. Specifically, we used a discrete-event simulation model of an assembly line as an example, which we created using the simulation software AnyLogic, as is shown in Figure 3. In our application, three parts *(A, B, C)* are assembled onto one main part, which is then delivered to the customer. To accomplish this, the main part and the elements to be mounted undergo preprocessing. The three parts to be assembled are collected in batches for preprocessing and then pass individually through a quality assurance station before reaching the assembly station. If the quality assurance process fails, then the elements must be reworked and preprocessed again. Each station, except for the preprocessing of the main parts, is operated by workers. In this process, quality assurance and reworking of parts A, B, and C are performed by workers from the same worker pool.

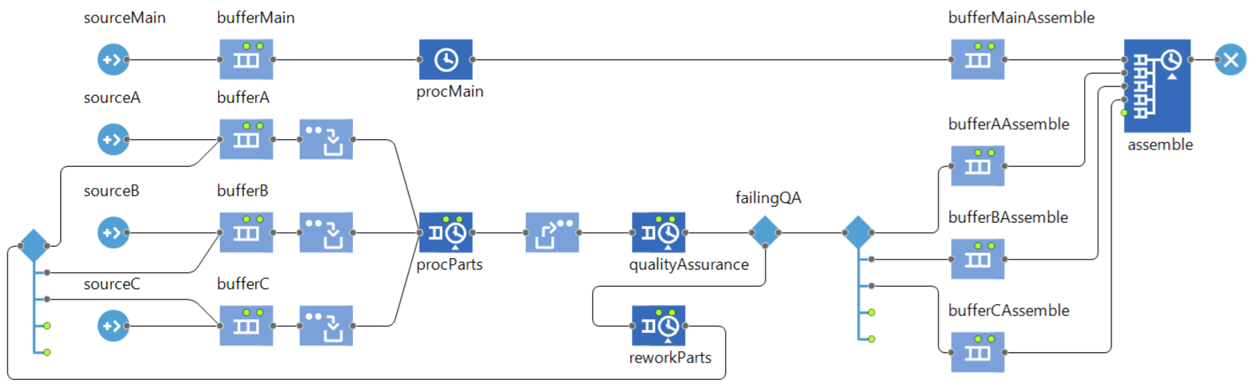

Figure 3: Screenshot of the assembly line in AnyLogic.

For the experiment design, the variable factors of the simulation model and their limits are summarized in Table 2. These factors mainly include time-specific factors (in minutes), such as the processing times and arrival rates of the various parts, as well as the assembly, quality assurance, and rework time. Other factors relate to the number of available workers for their respective station and the number of machines available for preprocessing the main part. These factors determine the number of possible processes that can run in parallel on the corresponding station. In relation to the buffers and batches of the system, the size of these can be varied.

Table 2: Variable factors and their limits of the simulation model.

Factor Name in the Model	Lower Limit	Upper Limit	Description
procTime{A,B,C}	1	10	Preprocessing time of element A, B, C
procTimeMain	1	10	Preprocessing time of the main product
batchSize{A,B,C}	1	15	Batch size for product A, B, C
procWorkerCount	5	15	Worker count for preprocessing
qaWorkerCount	10	20	Worker count for quality assurance
assembleWorkerCount	5	15	Worker count for the assembly station
arrivalRateMain	0.1	0.5	Arrival rate of the main product
arrivalRate{A,B,C}	0.1	1.5	Arrival rate of product A, B, C
assembleTime	1	5	Assembly time
bufferMainAssembleSize	100	300	Buffer size before the assembly station
buffer{A,B,C}AssembleSize	50	300	Buffer size before the assembly station
bufferProcPartsSize	50	300	Preprocessing A, B, C buffer size
bufferQASize	100	300	Quality assurance buffer size
bufferReworkSize	5	50	Rework station buffer size
qaTime{A,B,C}	0.2	1	Time required for quality assurance
reworkTime{A,B,C}	0.5	3	Time required for reworking
procMainCount	5	15	Number of machines for preprocessing

For each experiment, we measured the average utilization of workers, stations, and buffers across all replications. Other key metrics include the cycle time and average time in a queue of a part, as well as the maximum number of parts in the respective buffers during an experiment. The last metric is particularly interesting to determine which experiments have not only high buffer utilization but also reach their maximum capacity at least for a short time during a simulation run.

As a full factorial experimental design was not feasible due to the large number of possible combinations in our factor table, we used Latin hypercube sampling, which maintains good space-filling properties of our factor space (Viana 2013). For the Latin hypercube sampling, we used a correlation-minimizing sampling criterion to generate an orthogonal and correlation-free experimental design with ten thousand rows. After mapping the sampling results of the Latin hypercube to our factor table of the simulation model, this resulted in ten thousand experiments that were conducted in the data farming study. For each experiment, ten replications were performed, with each replication having a simulation run time of four days. The first day was not included in the output data due to the transient phase of the system. This resulted in a dataset with the size of ten thousand data points per factor, which is a rather small size for very complex models, but sufficient for our academic use case to demonstrate the visual problems.

4.2 How Not to and How to Visualize Your Simulation Output Data

In this chapter, we delve into the analysis and uncover common errors in visualizations making the interpretation and evaluation of the data difficult or misleading. To create these visualizations, the Python library Seaborn was used. It should be noted that due to the limited scope of this publication, we cannot cover all the methods, visualizations, and insights that were utilized and gained through an iterative process of adapting data mining methods and visualizations. Therefore, we present only selected results that highlight some common problems.

In a data farming study, the first questions to ask are: How are the output parameters distributed, and which ones are significant? To answer these questions, we analyze the output parameter distributions and identify any positive or negative anomalies. We then investigate what may have caused these anomalies and derive rules that can be applied to the system. Figure 4 displays the utilization of various machines and workers using a violin plot, which combines box whisker plots with density estimation.

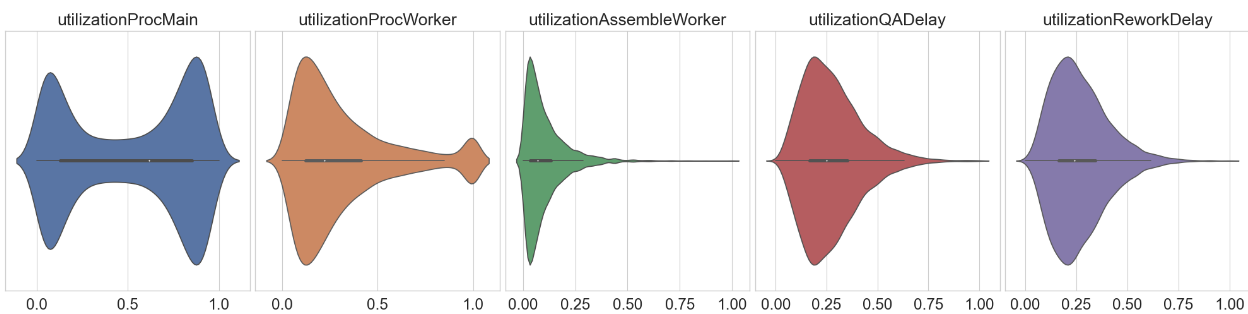

Figure 4: Problematic display of different utilizations of the simulation model with violin plots.

While density estimation can produce a less cluttered and more interpretable plot than a histogram, it may introduce distortions if the underlying distribution is bounded or not smooth. This problem is apparent here, and in fact the plots convey that there are utilizations, that are above 1 and below 0, which cannot be the case. This problem is apparent in Figure 4, where the plots suggest there are utilizations above 1 and below 0, which is not possible. In addition, statistical measures such as quantiles and means are not visible, x-axis labeling is not consistent, and large amounts of outliers in the distributions may be hidden. Consequently, this figure fails to meet the requirements for an expressive visualization. Therefore, a combination of visualizations and their corresponding benefits is often the best solution, provided they are consistent with one another.

In Figure 5, we demonstrate this approach by combining the advantages of histograms, density estimation, and box-whisker plots in a consistent representation, using a qualitative color palette to meet the requirements for the visual functions: identification, comparison, distinction. In particular, the histograms reveal that Figure 4 hides many experiments in which the preprocessing workers have a utilization close to 100%. After performing the statistical analysis on all output parameters, the most relevant outputs for this simulation model according to their variance and anomalies are the utilization of the preprocessing workers and the station for preprocessing the main part as well as the total throughput of the system.

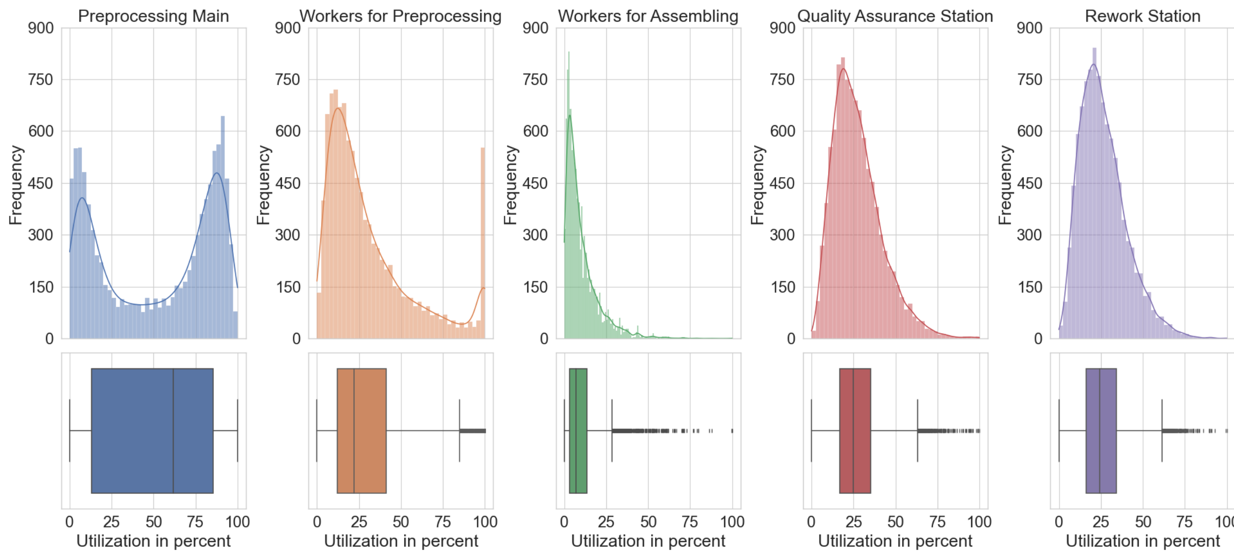

Figure 5: Appropriate visualization of the average utilizations in percent for different stations and workers using histograms and box-whisker plots.

If a distribution exhibits a very one-sided peak, and lacks significant variation or anomalies, it will be excluded from further analysis because they will not reveal interesting patterns. Next, we applied the k-means++ clustering algorithm to group experiments based on similarity of their relevant output performance measures. As a result, experiments within the same cluster show significant similarity with respect to the selected output performance measures, while experiments belonging to different clusters show significant dissimilarity. We have performed limited hyperparameter tuning, as this is not topic of this paper, but we achieved a good separation of the data with five clusters. Figure 6 presents the results of the clustering algorithm as a scatterplot matrix, where each data point represents one simulation experiment.

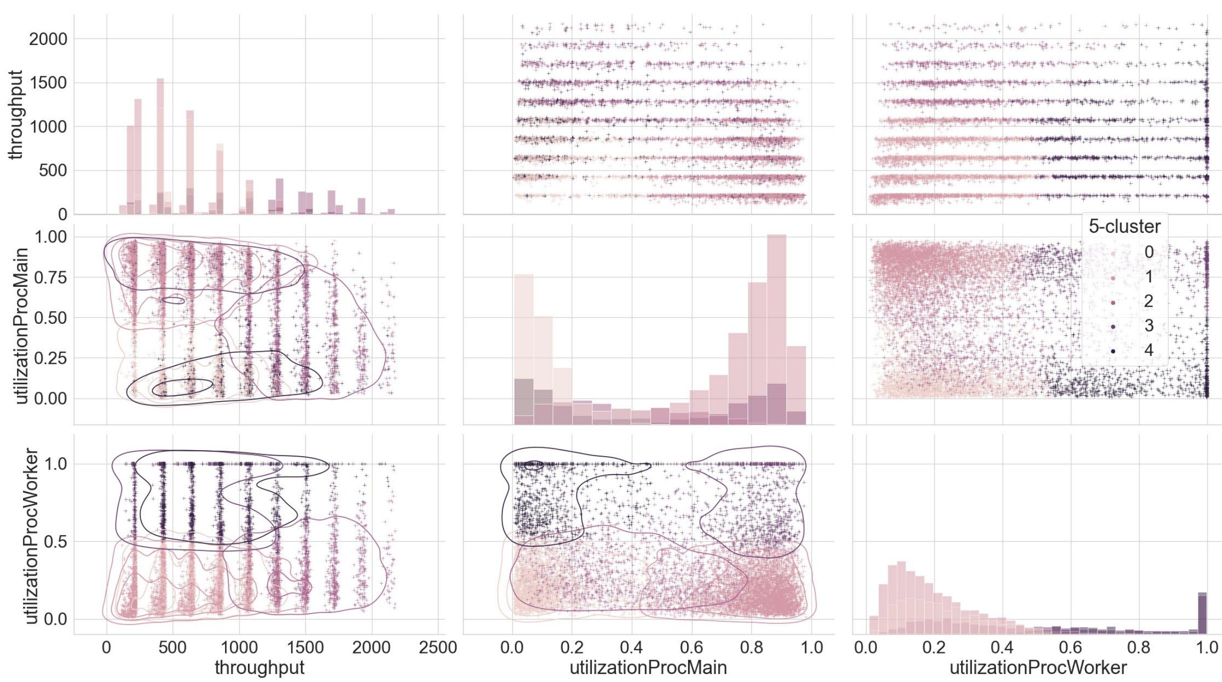

Figure 6: Inaccurate and obscure scatterplot matrix for the clustered and relevant output dimensions.

The representation of clustering results, as shown in the Figure 6, does not facilitate the identification of system behavior that can be classified as good or bad. This is mainly due to issues such as visual clutter, an incorrect color palette, information overload, and an unclear legend that obscures data points. The use of crosses to represent data points further contributes to the confusion in the scatterplot matrix. The attempt to depict clusters by using a sequential color palette and density estimation fails, making it impossible to differentiate and evaluate individual clusters in both the scatterplots and the histograms on the diagonal. In contrast, Figure 7 presents a more effective representation of multidimensional cluster analysis. It combines scatterplots, stacked histograms with equal numbers of bins, and abstract dense representations of the clusters using a qualitative color palette with proper labels. This visualization allows for the classification, location, and comparison of clustering results, revealing distinguishable patterns in the output dimensions that enable the analysis and grouping of system behavior in terms of qualitative behavior. Based on the results, we can identify Cluster 3 (green) as a group of experiments with good performance. It achieves the highest throughput while maintaining an average utilization of the main preprocessing station and the workers responsible for preprocessing parts. Conversely, Cluster 5 (blue) summarizes experiments with poor system performance.

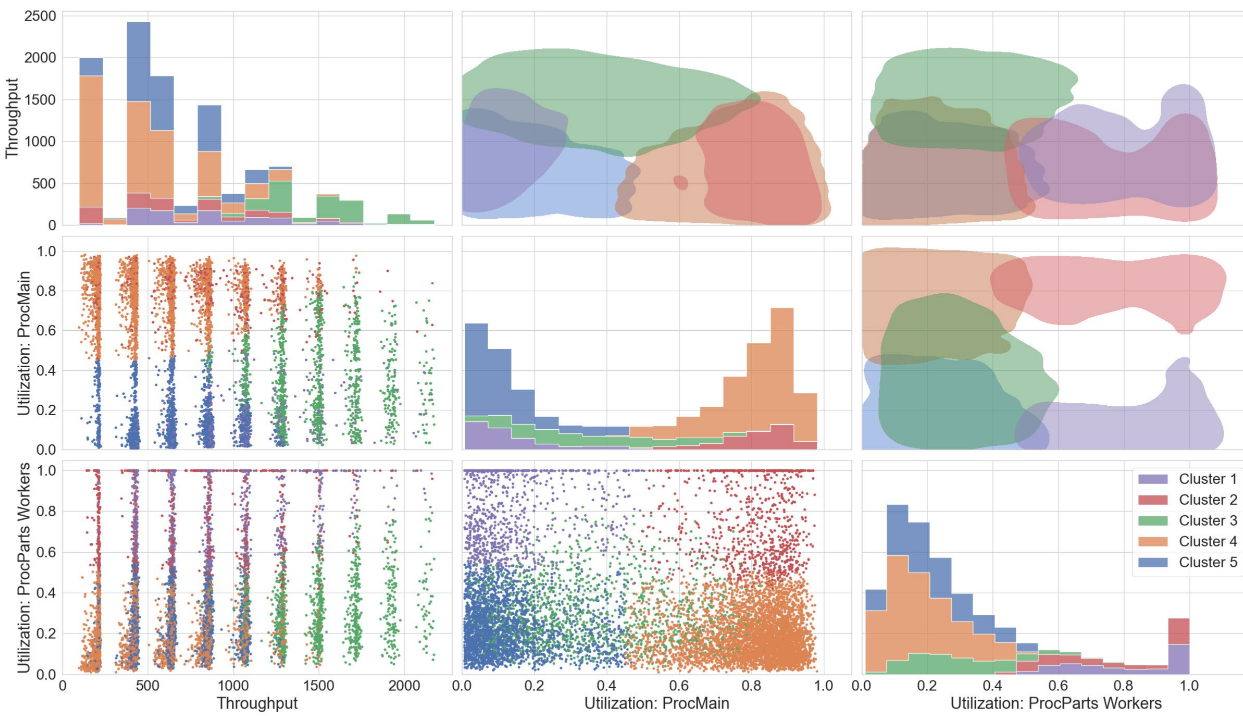

Figure 7: Scatterplot matrix with stacked histograms for the distribution of values on the diagonal and abstract density representations for the clusters in the upper right corner for the clustered outputs: throughput, preprocessing station utilization, and preprocessing worker utilization.

Our next objective is to identify the factor combinations that contribute to system performance. To achieve this, we can leverage supervised data mining techniques or skewed distribution analysis, as recommended by KDS. In Figure 8, we applied skewed distribution analysis to our good cluster by using the cluster as a filter to identify skewed factor value distributions. The greater the degree of skewness in a factor value distribution, the more significant and decisive that factor value becomes for the good cluster or experiment subset. Since our clusters are based on multidimensional patterns, we can determine the relationship between system behavior and its factors. But Figure 8's pie charts fail to show the distributions of arrival rates and hence do not display any skewness.

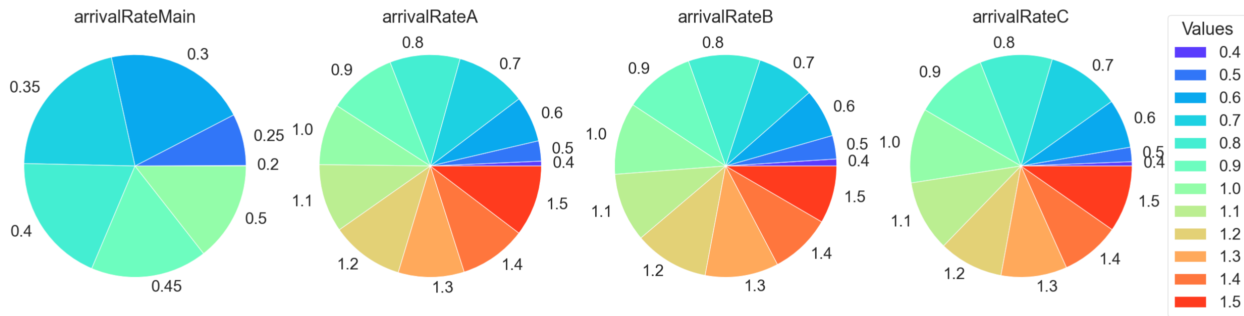

Figure 8: Ambiguous and confusing presentation of the arrival rates with pie charts.

The limitation arises from the fact that humans have a poor ability to perceive and process areas and angles accurately (Schwabish 2021). For this reason, it is advised to avoid pie charts and the use of areas and angles to represent data completely. This is compounded by using a rainbow color palette, which further obscures the data presentation by seemingly illuminating some values and by creating a false impression of some values being more significant than others, even when they are not. In this type of presentation, it is impossible to identify the factors with peculiarities or anomalies because they also have different scales. However, if we normalize all the factors before plotting them in a box-whisker plot, it immediately becomes apparent that the arrival rates, particularly those of the main part, have a substantially different distribution, as displayed in Figure 9. Normalization is necessary for this presentation because the different scales of the factors have no importance when considering them collectively. Other factors do not appear to influence cluster affiliation based on this analysis.

In summary, we can conclude that a high arrival rate is necessary to achieve good system performance. Although this finding may seem obvious for our simple simulation model, our study highlights how poor visualizations can conceal or hide valuable insights into the output data and the behavior of the system.

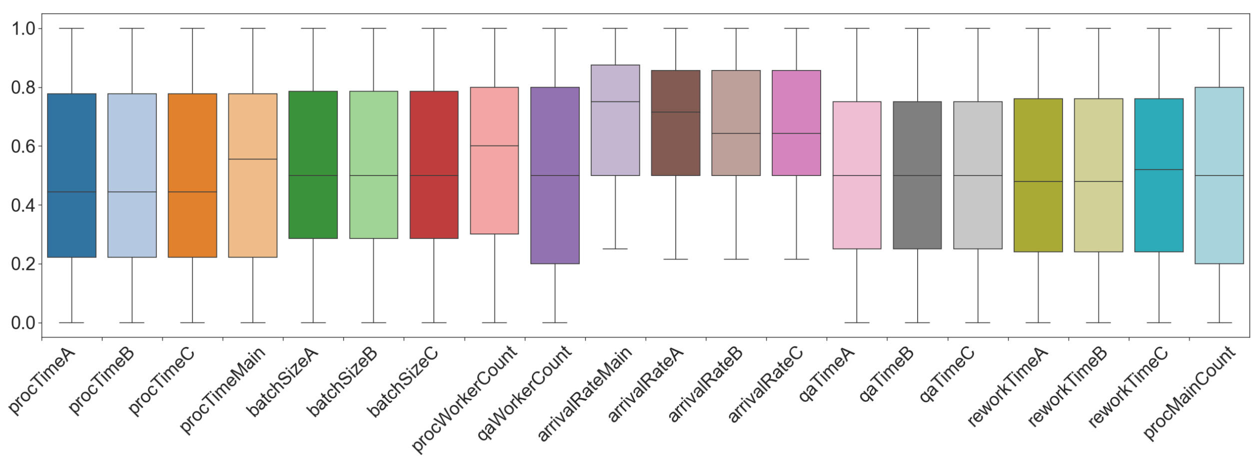

Figure 9: Box-whisker plot of most of the normalized factors filtered by the cluster with the good performance.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we have successfully shown why it is important to create visualizations with their visual function, and the perception of humans in mind as well as being consistent with the underlying data. To achieve this, we have discussed some fundamental requirements for visualizations that should always be considered, as there is a fine balance between information overload and presenting too little information to understand the data. We have successfully applied these principles to a data farming study, while highlighting how poor visualization choices and small mistakes can conceal information and obscure results. Throughout this approach, we explained the errors and mistakes that may occur when visualizing simulation output data and provided solutions to these issues while also contributing to narrowing the gap between simulation and visual analytics research.

Some aspects of visual analytics are beyond the scope of this paper; therefore, we emphasized the need for further work in the context of simulation, visual analytics, and visualization of research results as accurate communication of results through visualizations is crucial for scientific discourse. In addition, to increase awareness for the importance of presenting research findings accurately and clearly, further research needs to be conducted on the specific requirements for interactions and other aspects that affect human perception.

Future developments in visual analytics may also incorporate machine learning and deep learning techniques with interactive visualizations to enable users to explore and analyze data in more intuitive and efficient ways. However, a crucial research question that arises here is how to communicate iteratively and interactively generated results comprehensibly with other researchers.

REFERENCES

Brewer, C. A. 2016. *Designing Better Maps*: *A Guide for GIS Users.* 2nd ed. Redlands, California: Esri Press.

Cui, W. 2019. "Visual Analytics: A Comprehensive Overview". *IEEE Access* 7:81555–81573.

Ellis, G., D. Keim, F. Mansmann, and J. Kohlhammer. 2010. *Mastering the Information Age – Solving Problems with Visual Analytics*.

Endert, A., W. Ribarsky, C. Turkay, B. W. Wong, I. Nabney, I. D. Blanco, and F. Rossi. 2017. "The State of the Art in Integrating Machine Learning into Visual Analytics". *Computer Graphics Forum* 36(8):458–486.

Feldkamp, N. 2021. "Data Farming Output Analysis Using Explainable AI". In *Proceedings of the 2021 Winter Simulation Conference,* edited by S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Feldkamp, N., J. Genath, and S. Strassburger. 2022. "Explainable AI for Data Farming Output Analysis: A Use Case for Knowledge Generation Through Black-Box Classifiers". In *Proceedings of the 2022 Winter Simulation Conference,* edited by B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C. G. Corlu, L. H. Lee, E. P. Chew, T. Roeder, and P. Lendermann, 1152–1163. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Feldkamp, N., S. Bergmann, S. Straßburger, and T. Schulze. 2020. "Visualization and Interaction for Knowledge Discovery in Simulation Data". *53rd Hawaii International Conference on System Sciences* :1340–1349.

Feldkamp, N., S. Bergmann, and S. Strassburger. 2015. "Knowledge Discovery in Manufacturing Simulations". In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation,* edited by S. J. Taylor, N. Mustafee, and Y.-J. Son, 3–12. 3rd ed. New York, NY: ACM.

Ham, D.-H. 2010. "The State of the Art of Visual Analytics". In *Proceedings of the EU-Korea Conference on Science and Technology,* edited by J. H. Lee, H. Lee, and J.-S. Kim, 213–222. Berlin, Heidelberg: Springer Berlin Heidelberg.

Horne, G. E., and T. E. Meyer. 2005. "Data Farming: Discovering Surprise". In *Proceedings of the 2005 Winter Simulation Conference,* edited by M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 1082–1087. New York: Association for Computing Machinery.

Horne, G. E., and K.-P. Schwierz. 2008. "Data Farming around the World Overview". In *Proceedings of the 2008 Winter Simulation Conference,* edited by S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler, 1442–1447. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Keim, D., F. Mansmann, and J. Thomas. 2010. "Visual Analytics: How much Visualization and how much Analytics?". *ACM SIGKDD Explorations Newsletter* 11(2):5–8.

Keim, D., G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. 2008. "Visual Analytics: Definition, Process, and Challenges". In *Information Visualization*: *Human-Centered Issues and Perspectives,* edited by D. Hutchison, J.-D. Fekete, T. Kanade, A. Kerren, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, and C. North, 154–175. Berlin, Heidelberg: Springer.

Keim, D. A. 2001. "Visual Exploration of Large Data Sets". *Communications of the ACM* 44(8):38–44.

Kleijnen, J. P. C., S. M. Sanchez, T. W. Lucas, and T. M. Cioppa. 2005. "State-of-the-Art Review: A User's Guide to the Brave New World of Designing Simulation Experiments". *INFORMS Journal on Computing* 17(3):263–289.

Law, A. M. 2003. "How to Conduct a Successful Simulation Study". In *Proceedings of the 2003 Winter Simulation Conference,* edited by S. Chick, P. J. Sanchez, D. Ferrin, and D. J. Morrice, 66-70. New York: Association for Computing Machinery.

Lechler, T., M. Sjarov, and J. Franke. 2021. "Data Farming in Production Systems - A Review on Potentials, Challenges and Exemplary Applications". *Procedia CIRP* 96:230–235.

Lucas, T. W., W. D. Kelton, P. J. Sánchez, S. M. Sanchez, and B. L. Anderson. 2015. "Changing the Paradigm: Simulation, Now a Method of First Resort". *Naval Research Logistics (NRL)* 62(4):293–303.

Matković, K., D. Gračanin, and H. Hauser. 2018. "Visual Analytics for Simulation Ensembles". In *Proceedings of the 2018 Winter Simulation Conference*: *Simulation for a Noble Cause,* edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 321–335. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Moreland, K. 2009. "Diverging Color Maps for Scientific Visualization". In *Advances in Visual Computing,* edited by D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, G. Bebis, R. Boyle, B. Parvin, D. Koracin, Y. Kuno, J. Wang, R. Pajarola, P. Lindstrom, A. Hinkenjann, M. L. Encarnação, C. T. Silva, and D. Coming, 92–103. 5th ed. Berlin: Springer.

Munzner, T. 2014. *Visualization Analysis and Design.* 1st ed. Boca Raton, London, New York: CRC Press.

Mustafee, N., and J. H. Powell. 2018. "From Hybrid Simulation to Hybrid Systems Modelling". In *Proceedings of the 2018 Winter Simulation Conference*: *Simulation for a Noble Cause,* edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 1430–1439. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Sanchez, S. M. 2014. "Simulation Experiments: Better Data, Not Just Big Data". In *Proceedings of the 2014 Winter Simulation Conference,* edited by A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, 805–816. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Sanchez, S. M., and P. J. Sanchez. 2017. "Better Big Data via Data Farming Experiments". In *Advances in Modeling and Simulation,* edited by A. Tolk, J. Fowler, G. Shao, and E. Yücesan, 159–179. 1st ed. Springer International Publishing.

Schwabish, J. A. 2021. *Better Data Visualizations*. New York: Columbia University Press.

Spence, R. 2014. *Information Visualization: An Introduction.* 3rd ed. Cham: Springer.

Strassburger, S., S. Bergmann, N. Feldkamp, K. Sokoll, and M. Clausing. 2018. "Data Farming Research Project with Audi and VW". In *Plant Simulation Worldwide User Conference,* October 16th -18th, Stuttgart, Germany.

Tolk, A., A. Harper, and N. Mustafee. 2021. "Hybrid Models as Transdisciplinary Research Enablers". *European Journal of Operational Research* 291(3):1075–1090.

Viana, F. 2013. "Things You Wanted to Know about the Latin Hypercube Design and Were Afraid to Ask". In *10th World Congress on Structural and Multidisciplinary Optimization,* 19.-24.05, Orlando, Florida, USA, 1–9.

Ware, C. 2021. *Information Visualization: Perception for Design.* 4th ed. Cambridge, MA: Morgan Kaufmann.

AUTHOR BIOGRAPHIES

JONAS GENATH has been working since the beginning of 2020 as a research associate in the Group for Information Technology in Production and Logistics at the Ilmenau University of Technology. He received his M. Sc. from the Ilmenau University of Technology in 2019. His research topics focus on the automation of knowledge discovery in simulation data, data science and visual analytics. His email address is jonas.genath@tu-ilmenau.de.

STEFFEN STRASSBURGER is a professor at the Ilmenau University of Technology and head of the Group for Information Technology in Production and Logistics. Previously he was head of the "Virtual Development" department at the Fraunhofer Institute in Magdeburg, Germany and a researcher at the Daimler Chrysler Research Center in Ulm, Germany. He holds a Doctoral and a Diploma degree in Computer Science from the University of Magdeburg, Germany. His research interests include distributed simulation, automatic simulation model generation, and general interoperability topics within the digital factory and Industry 4.0 context. His email address is steffen.strassburger@tu-ilmenau.de.