

A systematic literature review of machine learning canvases

Lukas-Walter Thiée  ¹

Abstract: The use of machine learning technology is still significantly lower in small and medium sized enterprises than in large enterprises. It seems that there are specific challenges in the implementation of data-driven methods, that hinder SMEs in their adoption. One approach to support the initialization and execution of such methods is the use of boundary objects, e.g., canvases, serving as a visual communication document. As it is not clear which approaches are being pursued in detail and how they are interrelated, in this paper, a systematic literature review is being presented, that identifies and analyzes 18 canvas artifacts. These canvases represent the status quo and they can be grouped into four distinct categories of different foci. The aggregation of the fields and questions provides an essence of canvas contents, to point out gaps and ultimately to expand the canvas approach as well as ML adoption.

Keywords: Machine Learning, Canvas, Literature Review, SME

1 Introduction and Approach

The extent to which companies apply machine learning (ML) has increased in recent years and with over a quarter of German companies using multiple ML technologies today [ID20], it is far from a niche technology. While larger companies have been taking advantage of data-driven analytics for some time, mainly with a focus on process optimization [WG19], [VD18] and product development [WG18], for small and medium sized companies (SMEs) in Germany it is still difficult to start and use ML applications in their businesses [Mi20a]. As of 2020 only 10% of the small companies utilize multiple ML technologies and over 22% admit that ML applications are still not a topic in the company at all [ID20]. Furthermore, only 29% of SMEs assess ML as a driver for innovation and product development [ID20], which could indicate an imbalance in the assessment of opportunities and challenges of data science (DS), artificial intelligence (AI) in particular. Nevertheless, two out of five SMEs are beginning to plan digitalization projects and another 29% are considering to do so [Kf20]. Even though studies have shown benefits of data usage for SMEs [Mo18], SMEs have been focusing descriptive approaches and conventional business intelligence [BF18], rather than leveraging predictive or prescriptive data analytics to its full potential [Co16]. It seems that to the same extent as SMEs were challenged with the introduction of information technology in its beginnings [Gh11], they are now challenged to implement advanced data analytics and data strategies

¹ Leuphana Universität Lüneburg, Institut für Wirtschaftsinformatik, Universitätsallee 1, 21335 Lüneburg,
lukas-walter.thiee@leuphana.de,  <https://orcid.org/0000-0002-1998-376X>

[VLF15]. However, in practice there seem to be specific challenges for SMEs regarding the initialization and successful execution of ML projects [BvK20]. A regional survey in Germany identifies high investment cost, data collection and processing, shortage of specialists and AI maturity as the main barriers for companies [Mi20b]. One of the main reasons for the slow adoption of ML technology in SMEs is probably the lack of resources [Kf17]. Compared to larger companies, SMEs have less financial power and fewer or no skilled personnel for specific tasks in ML projects. Consequently, this gives rise to a variety of challenges. Lack of management focus and experience [Mo17], lack of internal and external experts [Ru15], or complexity of (sequential) data [FRL17] are exemplary challenges that can impede both project initiation and success. One of these challenges is the detailed description and expedient guidance in the use case definition and execution phase of an ML project. Naturally, these projects include multiple stakeholders at varying levels of data-literacy [Ke21], which unfortunately promotes the development of communication silos [ID20]. This communication barrier, internally and externally, must be overcome. Therefore, in addition to general methods that foster ML adoption, such as design thinking workshops, specific tools have been proposed, such as the DUCAR process model for smart picking of ML use cases [Sc20], or procedures to identify and prioritize use cases [Ku17], [Sc16]. Also the use of Guided Analytics [Bo19] has been proposed, so that non-experts can access ML tools [Qi18]. Nevertheless, this approach currently still mainly focuses model tuning. Another promising approach to facilitate ML communication and execution is the use of canvases. “A canvas is just a visual chart to describe a complex object, in a better way than a simple text document. [...] blocks are arranged on the chart in a way that makes visual sense” [Do19]. Canvases are a kind of boundary object, i.e., a document that serves as a communication platform between multiple stakeholders, motivates cross-disciplinary collaboration [KPD18], and provides common identity [SG89]. Originating from the Business Model Canvas by Osterwalder and Pigneur [OP10], this approach has been further developed and analyzed in many ways, for example in data-driven business models [Ha14]. However, it is not clear which of these canvas approaches is most promising for initial ML projects.

Research Approach and Contribution. For this reason, it is essential to review relevant resources systematically. Comparing different data-related canvases not only streamlines research in this field, but also provides useful guidance in business practice. A compilation of existing tools and a possible extension that can easily be applied in practice, can help bridging between the aforementioned corporate silos, i.e., decision-makers, developers, domain experts, and external partners, as well as promote the execution of ML tasks. Furthermore, the result of this review will be the foundation of conducting empirical research in longitudinal case studies to investigate the applicability and usefulness of such boundary objects. The detailed research questions for this paper are: (1) Which canvas models, that address ML or AI implementation, are available, and which contents do they cover?; and (2) Where are gaps and what are potential extensions of these canvases in order to address specific challenges and needs in initial ML projects? Providing answers to these questions might support initialization and successful execution of ML projects in small organizations. Therefore, the goal for this review is to find relevant canvas tools, compare these tools systematically, and potentially enhance them.

2 Methodology: Literature search process documentation

In order to find existing evidence, identify gaps, and build an appropriate background on the topic a systematic literature review is being conducted [Ki04]. The review shall identify relevant sources within a defined scope, synthesize the findings [WW02], and most of all provide guidance on research and practice [Sc15]. As rigor and reproducibility are key to qualitative IS reviews [Br09], the search process is documented in the following. The selection of databases (see Tab. 1) comprises ten international renowned databases, such as the Web of Science Core Collection, the ACM digital library and IEEE Xplore. The latter belong to the most impactful databases in computer science [KJZ16]. The selection of databases was inspired by the journal rating for IS literature from VHB JOURQUAL3 [HS15], which also comprises IS conference proceedings. The AIS Electronic Library features important IS outlets, such as MISQ or BISE. Furthermore, two online archives, Google Scholar and ArXiv.org; as well as three German online libraries were included, because they basically supplement the corpus of resources. To further integrate practitioner views, the Harvard Data Science Review, which is a non-peer-reviewed open access journal of the MIT Press, was also considered.

Database	Results	Dupl.	TAK	
AISeL	4	4	4	
EBSCOhost BSC	2	2	2	Backward:
ACM Digital Library	2	1	0	+15
Emerald Insight	0	0	0	
ScienceDirect	3	3	1	
SCOPUS	8	4	1	Forward:
IEEE Xplore	1	1	1	+1
Web of Science	15	14	2	
Google Scholar	3	0	0	
ArXiv	1	0	0	Language:
Taylor&Francis	3	1	0	-2
Springer Link	3	3	1	
HMD	0	0	0	
Duncker & Humblot	0	0	0	Artifact Identity:
RonPub	0	0	0	-1
Harvard DS Review	0	0	0	
Total	45	33	12	25

Tab. 1: Databases and literature search process results

The formulation of the search terms is very important in this matter, because there are often no clear dividing lines between the terms, e.g., the phrases “data mining”, “data-driven”, “digitalization”, or “big data analytics” all have a large intersection, which is mainly due to the fact that in the field of IS there is a plethora of terms and abbreviations that describe the topic area or sub-areas. Not even the distinction between AI, ML and deep learning is unambiguous. For example, an initial search in the EBSCOhost Business Source library, yields over 30,000 results. Therefore, we limit the search to the search string: *{ "Machine Learning Canvas" OR "Artificial Intelligence Canvas" OR "ML Canvas" OR "AI Canvas" }*. This initially excluded phrases such as digital canvas or data

science canvas. The query was individually adapted to the requirements of the respective database. No keyword was integrated that was specifically suitable for SMEs, since it was assumed that so far rather general approaches are available. The query was executed on all fields, i.e., title, abstract, keywords and full text, and the time span for results was limited to the years 2000 to 2021. The search was conducted end of March 2021. As the search term is rather strict, there were 45 initial hits within the included databases (see Tab. 1). Eliminating duplicates left 33 results and screening titles, abstracts, and keywords (TAK) reduced the results to 12 articles. The screening excluded hits that didn't relate to ML or AI canvases or projects, such as the "Business Ethics Canvas" [VHR20]. A backward search through the listed references as proposed by Webster and Watson [WW02] complements the search and adds 15 articles to the results. The backward search was performed in such a way that a TAK screening was performed for references that showed a promising title related to canvases, such as "Data-Driven Business Models" [BT19]. Two articles were excluded from the results, because they were not written in English and another article was excluded, because it was a foundational research article to a different article in the results, describing the same artifact, "ML-Process Canvas" [Zh19]. The results all originated from the time window of the years 2016-2021. This period is very recent, therefore, conducting a forward search in the mentioned databases didn't reveal any further relevant results. However, a sample forward search on Google yielded citations on the professional online network LinkedIn, which produced another result [Sc18]. In total there were 25 relevant articles left, which contained 18 canvases for in-depth analysis, which is the answer to part one of research question one (Tab. 2).

Year	Source	Canvas Artifact	#Fields	Structure
2016	[MK16]	Data Canvas: Data-Need Fit	10	F, D, E
2017	[SN17]	Data Value Map	14	F, S, D
2017	[He17]	Digitalization Canvas	9	F, S, Q
2018	[AGG18]	AI Canvas	7	F, Q
2018	[DR18]	AI Canvas	8	F, Q
2018	[Sc18]	The ML Canvas (Big Data MBA Version)	+2	F, D, Q
2018	[Sc18]	Hypothesis Development Canvas v1.1	10	F, Q
2019	[BT19]	Data Insight Generator	6	M, D, Q
2019	[Do19]	Machine Learning Canvas v0.4	10	F, D, Q
2019	[KM19]	Data Innovation Board	14	F, S, Q
2019	[KMK19]	Data Collection Map	12	F, E
2019	[Za19]	AI Project Canvas	9	F, S, Q
2020	[EL20]	AI performance canvas (prototype)	10	F, S, Q
2020	[FBP20]	Data Product Canvas	7	F, Q, E
2020	[HST20]	Key Activity Canvas	10	M, Q, E
2020	[KS20]	Canvas for the use of AI (author)	7	F, Q
2020	[Zh20]	ML Lifecycle Canvas	6	M, D, Q
2021	[Ke21]	Enterprise AI Canvas	12	F, S, Q

Notes: F = Fields, M = Matrix, S = Sections, E = Examples, D = Descriptions, Q = Questions

Tab. 2: Literature search results – 18 canvas artifacts

3 Analysis and Results

The analysis of the canvases is performed in two parts, structural analysis and content categorization. The structural analysis considers the layout and the process of filling. For this purpose, both the artifacts themselves and a full-text analysis are performed. The process in the full-text analysis followed in this qualitative content analysis is based on inductive formulation of categories [Ma10], that were given by the captions in the canvases. Also, it is being determined whether and which roles or persons are being addressed to fill the respective canvas and if the filling is self-sufficient and detailed guidelines published. Further, if applicable, the scientific derivation and evaluation of the artifacts are investigated.

Structural analysis. Almost all artifacts are structured as a canvas with fields that have a descriptive title accompanied by either questions or examples (see Tab. 2, Structure column). A few use a 2-d matrix arrangement with rows and columns. Some artifacts provide sections or arrows in the layout to guide the user in filling in the fields, e.g., the “Key Activity Canvas” provides dotted arrows to visualize interactions between the customer, the company, and the partners [HST20], or Zawadzki draws arrows that indicate how to go from section to section in the “AI Project Canvas” [Za19]. 4 articles provide written guidelines how to proceed in filling in the canvas fields, such as “Explore, Ideate, Evaluate” [KM19], “Think, Validate, Know” [BT19], design loops [Zh20], or agile development [He17]. Whereas 3 others point to the iterative nature of the tool [Do19], [Sc18], [HST20], the remaining 11 do not specify the process (once or iterative), so that a single pass of the filling has to be assumed. In terms of size, or number of fields respectively, the 18 artifacts range from 7 fields in the “AI Canvas” [AGG18] to 14 fields in the “Data Innovation Board” [KM19]. The highest number of cells, results from one of the matrix approaches, namely 21 cells in the “Key Activity Canvas” [HST20].

As an overarching finding, it can be noted that all articles see data-driven projects, and the respective canvas in particular, as an interdisciplinary task. Nevertheless, 5 of the 18 articles don’t specify the person or department which should fill in the canvas. 5 others only mention general terms, such as “business stakeholders” [Sc18], “heterogeneous stakeholder groups” [BT19], “pioneers” [DR18], or “different departments and diverse expertise” [MK16]. Specifically mentioned are Data and AI project managers, IT departments, domain experts, service design teams, data science teams, “senior executives, middle management, frontline staff, business stakeholders, technology stakeholders and customers” [SN17], or in general “managers, who provide the glue between everyone” [Do19]. Regarding the scientific derivation, 6 articles explicitly name Design Science Research, Action Design Research, Research through Design, or Design Thinking as methodological procedures in their articles, e.g., applying questionnaires, triangulation, or design principles as methods. 5 of these also describe the evaluation procedure in their research, which is predominantly a focus group workshop. Also, the ontology, which the artifact is based upon, is mentioned in 2 articles, namely for the “Data Collection Map” [KMK19], and for the “Data Innovation Board” [KM19]. 2 other articles describe the interviews and workshops conducted as part of a case study [He17], [MK16]. [DR18] and

[Do19] ground their artifacts on the Business Model Canvas by [OP10]. However, nearly half of the results (8/18) don't specify the scientific method, which the artifact is based upon, and these contributions also don't mention the evaluation technique applied.

Categorization. In the second part of the analysis, the individual canvases and their core contents and objectives are being investigated. 6 of the 18 canvases explicitly label the canvas with the term AI, e.g., the “AI Canvas” by Agrawal et al. (2018), and 3 assign ML to their artifact, e.g., the “ML Lifecycle Canvas” by Zhou et al. (2020). The naming already indicates that different foci are being set. As the artifacts all belong to the same realm of data science, categorizing them is ambitious due to the proximity of their contents. Nevertheless, four categories can be proposed, as summarized in Fig. 1. This methodological procedure is in line with Webster and Watson's (2002) call for a concept-centric approach in IS literature reviews [WW02].

ML / DS Focus	(AI) Project Focus	Data Value Focus	Data Source Focus
<ul style="list-style-type: none"> • [Do19], [Sc18], [Sc18], [Zh20], [HST20] 	<ul style="list-style-type: none"> • [AGG18], [DR18], [Za19], [KS20], [Ke21], [EL20], [He17] 	<ul style="list-style-type: none"> • [FBP20], [SN17], [MK16], [KM19] 	<ul style="list-style-type: none"> • [BT19], [KMK19]

Fig. 1: Categories of canvases with different thematic foci

Machine Learning / Data Science Focus. The first category is formed by canvases with a technical focus on machine learning and data science. It includes the “Machine Learning Canvas v0.4” [Do19] and its extension the “Machine Learning Canvas (Big Data MBA Version)” [Sc18], as well as the “Hypothesis Development Canvas” [Sc18], the “ML Lifecycle Canvas” [Zh20], and the “Key Activity Canvas” [HST20]. These artifacts belong to this category, because they describe concrete machine learning or data science process steps, e.g., the definition of inputs and outputs or engineering of corresponding features. “The Machine Learning Canvas [is] the first step towards making sure you connect what ML can do to your organization’s objectives, and towards assessing feasibility. It should be filled in before starting any implementation work, and even before Exploratory Data Analysis” [Do19]. This artifact is the most technical, as it contains fields like “ML task” or “Features”. Dorard also integrates metrics – offline and online – and value proposition into his canvas. However, it is supposed to be an initial document, and “the canvas results [have to be] translated to a technical specification document” [Ma19] later on. Schmarzo builds upon Dorard’s canvas and proposes two additional fields, namely Prescription and Automation in the “Big Data MBA Version”, to adopt the canvas to “data science requirements” [Sc18]. This extension shall form the path from a small ML use case to an integrated, scaled application. Nevertheless, it’s not obvious how to really do that. Schmarzo also proposes the “Hypothesis Development Canvas v1.1” in order to facilitate „collaboration between the business stakeholder and the data science team to identify the hypothesis requirements that underpin data science engagement success” [Sc18]. Another rather technical approach is the “ML Lifecycle Canvas”. It’s a “conceptual design tool featuring the holistic visualization of cooperation among ML, users, and scenarios during the ML lifecycle” [Zh20]. This canvas is unique in the regard,

that it provides a detailed question list and “persona cards” to fill in the canvas. Questions like “Is there any ML model feasible for completing the required tasks?” illustrate the level of detail regarding the final ML design process. The authors align the questionnaire with questions from “existing guidebooks on human-AI interaction” [Zh20], like Google PAIR [Go19] or [Am19]. The “Key Activity Canvas” [HST20] is a matrix arrangement and integrates three views, Customer, Company and Partners, for the “methodological assistance (key activities) guiding the actual conceptualization of necessary activities in analytics-based services”.

(AI) Project Focus. The second category includes canvases with a holistic view on AI projects, “AI Canvas” by [AGG18] and the eponymous by [DR18], “AI Project Canvas” [Za19], “Canvas for the use of AI” [KS20], “Enterprise AI Canvas” [Ke21], “AI performance canvas” [EL20], and “Digitalization Canvas” [He17]. In principle, these could synonymously be applied to ML projects. However, the distinction from the first category arises from the fact that a focus is placed on the overall project rather than on the technical details, e.g., the cost and revenue structure of an AI project. The “AI Canvas” [AGG18] is supposed to support corporate decision-making through prediction models. “How can you decide whether employing a prediction machine will improve matters? The AI Canvas is a simple tool that helps you organize what you need to know into seven categories [Prediction, Judgment, Action, Outcome, Input, Training, Feedback] in order to systematically make that assessment” [AGG18]. It’s rather designed “for a non-technical audience” [EL20]. Technical details and business integration are not being covered [Ke21]. This is also true for the “AI Canvas” by Dewalt and Rands. They try to connect a business opportunity via a strategy with a solution using AI models. It’s the pathway to “Become an AI Company in 90 Days” [DR18]. Similarly, based on the original Business Model Canvas the “AI Project Canvas” by Zawadzki is helpful for “project managers” [Ke21] and is intended to “structure and convey the holistic idea of your AI project to others” [Za19]. Another way to “determine the relevance of Artificial Intelligence for your company” [KS20] is the “Canvas for the use of AI”, which is embedded in a corporate transformation process towards AI. The approach by Kerzel “Enterprise AI Canvas” [Ke21] features two parts, one with a technical focus and one with a business focus. It’s supposed to “bring business and data science experts together and systematically evaluate potentially new business opportunities” [Ke21]. Although part 2 “model and data view” shows similarities to Dorard's approach, technical details on modeling are omitted. The AI Project Focus category also includes the “AI performance canvas”, whose main objective is the “collaborative construction of performance goals for data & AI products in organizations” [EL20]. In this approach, there is a strong focus on feasibility of the product and legal/compliance issues in data governance. Nevertheless, only a prototype canvas is being presented and the “trigger questions” of the fields are not yet published. The “Digitalization Canvas” [He17] is the only canvas specifically designed for SMEs and corresponds to a holistic digitalization strategy. The approach promotes easy-to-implement and strategic data projects. While the canvas itself isn’t self-sufficient, there are detailed questionnaires within the case study. “The Digitalization Canvas together with a project portfolio defines a concrete roadmap to digitalization and summarizes the arguments to apply for the necessary budget” [He17].

Data Value Focus. The third category are canvases with a data value or data product focus, namely the “Data Product Canvas” [FBP20], the “Data Value Map” [SN17], the “Data Canvas – Data-Need Fit” [MK16], and the “Data Innovation Board” [KM19]. These canvases seek to identify and implement the value of data through tangible customer or user benefits and information gain through data. The “Data Product Canvas” [FBP20] can be used when an organization in an initiation phase aims to develop ideas for a data product. Similarly, the “Data Canvas and Data-Need Fit are intended to spark a discussion on available data in organizations among diverse stakeholders. The Data Canvas provides trigger questions and a visual representation that help to develop a common understanding of available data” [MK16]. “A Data-Need Fit is found when data sources contribute gain creators and pain relievers that users find valuable” [MK16]. “To facilitate a shared understanding for data initiatives” [SN17] is also the main objective of the “Data Value Map”. This canvas focuses the process from data creator to data user and emphasizes data governance topics, such as data principles and access, and business related topics such as cost reduction or revenue generation. Finally, the “Data Innovation Board” [KM19] features three design thinking steps, namely Exploration, Ideation, and Evaluation, and promotes the description of performance goals in “a visual collaboration tool that anyone can work with”. The artifact is clearly supposed to facilitate initial progress in a data-driven project, and the authors note that “it is a beginner’s tool [, which needs] to be accompanied by other visual tools with more specific views on technology and algorithms” [KM19].

Data Source Focus. The fourth category includes canvases that focus on explaining the data source and the data processing, in order to gain a better understanding of data as an asset: the “Data Insight Generator” [BT19] and the “Data Collection Map” [KMK19]. The “Data Insight Generator” [BT19] is a workshop canvas that connects key data resources with a value proposition for data-driven business models. It contains columns for Pipes, Analytics and Insight. The process of filling in is guided by the rows Think, Validate, and Know, which are to be processed one after the other. Finally, the “Data Collection Map was designed as an entry point in the ideation process of data-driven use cases. Hence, the purpose of the tool is to get people to think about data (e.g. clicks and engagement metrics) instead of IT systems (e.g. Google Analytics) and to raise the necessary data awareness about the available data resources within the organization” [KMK19]. It’s basically an add-on to the “Data Innovation Board” by Kronsbein and Mueller (2019).

Categorization of the fields and questions. As it is not sufficient to only categorize the artifacts on a title level, the fields are being analyzed. In order to conceptualize the core content of the canvases and thus answer the second part of research question one, all fields (or headers in matrix patterns) are being captured and groups as well as top categories are being proposed. This should clarify where the canvases overlap and where the focus has been placed so far. In total there are 163 fields in the results, e.g., the “Machine Learning Canvas v0.4” [Do19] contains 10 fields that contain one or more questions to guide the filling (see Tab. 2). Logically, this count includes multiple entries. Therefore, the fields that cover a similar area in terms of content are grouped together. In the next step, multiple entries are eliminated and fields are combined that either describe exactly or almost the

same thing. For example, fields like “Data Sources”, “AI data base”, “Metadata”, and all the various data types from the “Data Collection Map” [KMK19], belong to the same group of “Data Sources” (see Tab. 3). As most fields contain more than one guiding question or example and not every field headline describes the same content, in order to refine the assignment, all guiding questions or examples are also individually examined. Ultimately, 39 groups and 11 top categories can be created. The content intersection of the final clusters compiles the name of their top category, e.g., the groups “Data Quality”, “Data policies”, and “Data lifecycle” build the top category “Data Governance”. The result of this assignment is presented in Tab. 3.

Cat.	Group	#	Example Question
Business & Value	Strategy	13	“What trends, market facts are relevant for the topic [...]?” [KM19]
	Risks	10	“What risks are associated with the use of AI for our industry?” [KS20]
	Operat. value	7	“How does the use-case generate value?” [Ke21]
	Revenue	6	“How will the project generate revenue?” [Za19]
	Cost	9	“Will the project reduce internal costs [...]?” [Za19]
Product & Customer	Product/Service	14	“Which potentials in the production area can be leveraged?” [KS20]
	Delivery	7	“In which form do we provide the data service to our users [...]?” [FBP20]
	Customer	16	“Who is our customer?” [FBP20]
	Gains/Value	9	“What is the customer value of the hypothesis?” [Sc18]
	Pains/Needs	9	“What customer pain is the AI project solving?” [Za19]
Organi- zation	Implementation	8	“How might we implement the idea?” [KM19]
	Internal skills	7	“Is the required know-how for the implementation available inhouse?” [He17]
	Stakeholders	3	“Sponsor: Which senior manager is responsible?” [Ke21]
	Domain	2	“Which domain expertise is needed?” [Ke21]
	Partners	8	“Which external services and products are required?” [He17]
Tech nol.	Systems	7	“Which systems are required and already available to handle data?” [Ke21]
	Infrastructure	7	“How are the models served? Edge, on-premise or Cloud?” [Ke21]
	Integration	4	“Which networks along the value chain are necessary?” [KS20]
Data charact.	Data types	25	“What kind of data do we need for training?” [Zh20]
	Data sources	9	“Which raw data sources can we use (internal and external)?” [Do19]
	Data availability	14	“What data is currently collected in the organization?” [KM19]
	Data collection	8	“How might we collect the needed data?” [KM19]
	Data pipelines	3	“Which interfaces can I use to combine this data?” [BT19]
Data gov.	Data quality	6	“How is the validity of the data, [...] consistency, and completeness?” [He17]
	Data policies	9	“Are there any compliance requirements [...]?” [He17]
	Data lifecycle	5	“Determining the definition, [...] and retirement of data.” [SN17]
Pre- proc.	Data preparation	4	“What do we have to do to prepare the data [...]?” [HST20]
	Features	2	“Which features are likely important?” [Ke21]
	Inputs	4	“What are the model inputs?” [DR18]
Modeling	Learning	6	“Is there any ML model suitable for the available dataset?” [Zh20]
	Analytics	5	“With which data analytics methods do we generate insights [...]?” [FBP20]
	Interpretation	4	“How can we interpret the mined patterns?” [HST20]
	Prediction	5	“What should be predicted?” [Ke21]
	Decision	9	“Prescription: Once we have a prediction, what do we do?” [Sc18]
Evaluation	KPI Model	3	“Which key metric are you optimizing for?” [Za19]
	KPI Business	6	“Outcome: What are your metrics for task success?” [AGG18]
	Improvements	3	“How can you use the outcomes to improve the algorithm?” [AGG18]
	Automation	6	“When do we create/update models with new training data?” [Do19]
	Live / Ex-post	6	“Methods and metrics to evaluate the system after deployment [...].” [Do19]

Tab. 3 Categories and groups of all canvas fields and questions with examples

4 Discussion and Conclusion

The primary contribution of this review is the categorization of existing canvases. The review generates overview for initial ML projects. It complements project management approaches in software development, such as the v-model or scrum. Using one of the categorized canvases is a solid starting point. Practitioners can use the four categories as guidance and pick one of the mentioned canvases, e.g., if they want to explore their data, the canvases from the category "Data sources" will help. Although "an over emphasis on technology" [SN17] has been mentioned as a potential barrier, we feel that diving deeper into the technical details, i.e., data processing and modeling, is key to foster ML adoption. Therefore, with regard to research question two, the canvases with a ML/DS Focus are most suitable. SMEs can use the canvases and/or the catalog of questions to promote cooperation, e.g., with research or consultancy. The canvases are all standalone artifacts for valid use cases. Nevertheless, there is still room for improvement for three reasons. The first reason is, that there are still gaps regarding content. E.g., no detailed questions for hyper parameter tuning, visualization of results, or concept drift [We19] could be identified. Also scalability and feasibility checks were not mentioned. The second reason regards the applicability of the canvases. Clear guidelines on who, when, and how to use the canvas are needed, describing the explicit benefit. Otherwise, their usefulness is mitigated and scientific artifacts will not be favored against gray literature. A balance between detailed description, e.g., as in Google PAIR [Go19], and self-sufficiency, will provide the greatest benefit. Eventually, the evaluation benchmark of the canvases must be: "What's the artifact from the artifact?" Meaning that using the canvas facilitated building an ML application that provides value. The third reason is concerning the fact that "canvases help us ask the right questions, but they don't provide the answers" [DR18]. In order to lift the canvas approach from ideation to application, providing answers to the questions is necessary to foster initial ML adoption. E.g., guiding the software tool chain or the model selection as in the scikit learn cheat sheet² might be fruitful. Other potential extensions of the canvas approach could be the integration of cloud service platforms, e.g., AWS, or other MLaaS providers [RGC15], or questions regarding "ground truth" [SB14]. Future research may include three key points. First, the question list has to be compared to a) the challenges of ML adoption in SMEs and b) to existing process descriptions, such as CRISP-DM [WH00]. Second, a comprehensive canvas for initial ML projects can be conceptualized from the findings of this review. However, the level of detail of this canvas shouldn't be sacrificed for generalizability, which could be prevented by a layered approach. And third, this concept can then be used and evaluated in empirical research, especially in workshops and case studies. The conceptualization of a new artifact and its evaluation would also address the inherent limitations of this review paper, as personal bias and experience could not totally be omitted, especially in the categorization parts. Qualitative research through interviews regarding ideation, communication, and problem-solving could build the basis to assess the added and perceived value of the canvases.

² https://scikit-learn.org/stable/tutorial/machine_learning_map/index.html

Bibliography

- [AGG18] Agrawal, A.; Goldfarb, A.; Gans, J.: A Simple Tool to Start Making Decisions with the Help of AI. <https://hbr.org/2018/04/a-simple-tool-to-start-making-decisions-with-the-help-of-ai>, accessed 16 Apr 2021.
- [Am19] Amershi, S. et al.: Guidelines for Human-AI Interaction. In (Brewster, S. et al. Eds.): CHI 2019. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, May 4-9, 2019, Glasgow, Scotland UK. ACM, New York, NY, pp. 1–13, 2019.
- [BF18] Becker, W.; Feichtinger, C.: Data Analytics in jungen Unternehmen. Ergebnisse einer Online-Befragung. *Der Betriebswirt* 2/59, pp. 26–31, 2018.
- [Bo19] Bourcetet, A. et al.: Guided Machine Learning for Business Users. BLED 2019 Proceedings, 2019.
- [Br09] Brocke, J. vom et al.: Reconstructing the giant: on the importance of rigour in documenting the literature search process: Proceedings of the 17th European Conference on Information Systems, 2009.
- [BT19] Babett Kühne; Tilo Böhmman: Data-Driven Business Models - Building the Bridge between Data and Value. Twenty-Seventh European Conference on Information Systems (ECIS2019), Stockholm-Uppsala, Sweden., 2019.
- [BvK20] Bauer, M.; van Dinther, C.; Kiefer, D.: Machine Learning in SME: An Empirical Study on Enablers and Success Factors. AMCIS 2020 Proceedings, 2020.
- [Co16] Coleman, S. et al.: How Can SMEs Benefit from Big Data? Challenges and a Path Forward. *Quality and Reliability Engineering International* 6/32, pp. 2151–2164, 2016.
- [Do19] Dorard, L.: The Machine Learning Canvas. A handbook for innovators and visionary managers striving to design tomorrow's Machine Learning systems. DRAFT version 0.1 released on 12 February 2019, 2019.
- [DR18] Dewalt, K.; Rands, R.: Become an AI Company in 90 Days: The No-Bullshit Guide for Understanding AI, Identifying Opportunities, and Launching Your First Product. Prolego, 2018.
- [EL20] Engel, M.; Lang, F.: A Pilot Study on Designing a Data & AI Performance Canvas. AMCIS 2020 Proceedings, 2020.
- [FBP20] Fruhwirth, M.; Breidfuss, G.; Pammer-Schindler, V.: The Data Product Canvas - A Visual Collaborative Tool for Designing Data-Driven Business Models. BLED 2020 Proceedings, 2020.
- [FRL17] Funk, B.; Rettenmeier, M.; Lang, T.: Deep Learning auf sequenziellen Daten als Grundlage unternehmerischer Entscheidungen. *Wirtschaftsinformatik & Management* 5/9, pp. 16–25, 2017.
- [Gh11] Ghobakhloo, M. et al.: Information Technology Adoption in Small and Medium-sized Enterprises: An Appraisal of Two Decades Literature. *Interdisciplinary Journal of Research Business* 7/1, pp. 53–80, 2011.
- [Go19] Google PAIR: People + AI Guidebook. <https://pair.withgoogle.com/guidebook/>,

accessed 16 Apr 2021.

- [Ha14] Hartmann, P. et al.: Big Data for Big Business? A Taxonomy of Data-driven Business Models used by Start-up Firms. Cambridge Service Alliance, 2014.
- [He17] Heberle, A. et al.: Digitalization Canvas - Towards Identifying Digitalization Use Cases and Projects. *Journal of Universal Computer Science* 11/23, pp. 1070–1097, 2017.
- [HS15] Hennig-Thurau, P. D. T.; Sattler, P. D. H.: VHB-JOURQUAL3. <https://vhbonline.org/vhb4you/vhb-jourqual/vhb-jourqual-3/tabellen-zum-download>.
- [HST20] Hunke, F.; Seebacher, S.; Thomsen, H.: Please Tell Me What to Do – Towards a Guided Orchestration of Key Activities in Data-Rich Service Systems. In (Hofmann, S.; Müller, O.; Rossi, M. Eds.): *Designing for Digital Transformation. Co-Creating Services with Citizens and Industry. 15th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2020, Kristiansand, Norway, December 2–4, 2020, Proceedings*. Springer International Publishing; Imprint: Springer, Cham, pp. 426–437, 2020.
- [ID20] IDG Business Media GmbH: Studie Machine Learning 2020. <https://www.lufthansa-industry-solutions.com/de-de/newsroom-downloads/news/maschinelles-lernen-wird-in-deutschen-unternehmen-standard/>, accessed 4 Feb 2021.
- [Ke21] Kerzel, U.: Enterprise AI Canvas Integrating Artificial Intelligence into Business. *Applied Artificial Intelligence* 1/35, pp. 1–12, 2021.
- [Kf17] KfW Bankengruppe: Digitalisierung der Wirtschaft: breite Basis, vielfältige Hemmnisse. Unternehmensbefragung. <https://www.kfw.de/PDF/Download-Center/Konzernthemen/Research/PDF-Dokumente-Unternehmensbefragung/Unternehmensbefragung-2017-%E2%80%93-Digitalisierung.pdf>, accessed 7 May 2021.
- [Kf20] KfW Bankengruppe: Unternehmensbefragung 2020. Anteil der Digitalisierungsplaner stagniert auf hohem Niveau. <https://www.kfw.de/PDF/Download-Center/Konzernthemen/Research/PDF-Dokumente-Unternehmensbefragung/Unternehmensbefragung-2020-%E2%80%93-Digitalisierung.pdf>, accessed 7 May 2021.
- [Ki04] Kitchenham, B.: Procedures for Performing Systematic Reviews. Keele University Technical Report TR/SE-0401, 2004.
- [KJZ16] Kehrer, S.; Jugel, D.; Zimmermann, A.: A systematic literature review of big data literature for EA evolution. *Digital Enterprise Computing 2016, Lecture Notes in Informatics (LNI)*, 2016.
- [KM19] Kronsbein, T.; Mueller, R.: Data Thinking: A Canvas for Data-Driven Ideation Workshops. *Hawaii International Conference on System Sciences 2019 (HICSS-52)*, 2019.
- [KMK19] Kayser, L.; Mueller, R.; Kronsbein, T.: Data Collection Map: A Canvas for Shared Data Awareness in Data-Driven Innovation Projects. *Proceedings of the 2019 Pre-*

ICIS SIGDSA Symposium, 2019.

- [KPD18] Kollwitz, C.; Perez Mengual, M.; Dinter, B.: Cross-Disciplinary Collaboration for Designing Data-Driven Products and Services. Pre-ICIS SIGDSA Symposium on Decision Analytics Connecting People, Data & Things, San Francisco 2018, 2018.
- [KS20] Kreutzer, R. T.; Sirrenberg, M.: AI Challenge - How Artificial Intelligence Can Be Anchored in a Company. In (Kreutzer, R. T.; Sirrenberg, M. Eds.): Understanding Artificial Intelligence. Fundamentals, Use Cases and Methods for a Corporate AI Journey. Springer International Publishing; Imprint: Springer, Cham, pp. 235–273, 2020.
- [Ku17] Kuschicke, F. et al.: A Data-based Method for Industrial Big Data Project Prioritization. Proceedings of the International Conference on Big Data and Internet of Thing, pp. 6–10, 2017.
- [Ma10] Mayring, P.: Qualitative Inhaltsanalyse. In (Mey, G.; Mruck, K. Eds.): Handbuch qualitative Forschung in der Psychologie. VS Verlag, Wiesbaden, pp. 601–613, 2010.
- [Ma19] Marin, I.: Data Science and Development Team Remote Communication: the use of the Machine Learning Canvas: 2019 ACM/IEEE 14th International Conference on Global Software Engineering (ICGSE 2019). Montreal, Quebec, Canada, 25-26 May 2019. IEEE, Piscataway, NJ, pp. 18–21, 2019.
- [Mi20a] Mittelstand 4.0-Kompetenzzentrum Dortmund: Digital in NRW - Kompetenz für den Mittelstand. Künstliche Intelligenz im Mittelstand - Potenziale und Anwendungsbeispiele, 2020.
- [Mi20b] Ministerium für Wirtschaft, Arbeit und Wohnungsbau Baden-Württemberg: MONITORING-REPORT. Wirtschaft DIGITAL Baden-Württemberg 2020. https://www.wirtschaft-digital-bw.de/fileadmin/media/Dokumente/Studien/Monitoring_BaWue_2020_Langfassung.pdf.
- [MK16] Mathis, K.; Köbler, F.: Data-Need Fit – Towards Data-Driven Business Model Innovation. ServDes. 2016, Fifth Service Design and Innovation conference, 2016.
- [Mo17] Moeuf, A., Lamouri, S., Pellerin, R., Eburdy, R., Tamayo, S.: Industry 4.0 and the SME: a technology-focused review of the empirical literature. Proceedings of 7th International Conference on Industrial Engineering and Systems Management IESM, 2017.
- [Mo18] Mohd Selamat, S. A. et al.: Big data analytics-A review of data-mining models for small and medium enterprises in the transportation sector. WIREs (Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery) 3, 2018.
- [OP10] Osterwalder, A.; Pigneur, Y.: Business model generation. A handbook for visionaries, game changers, and challengers. Wiley&Sons, New York, 2010.
- [Qi18] Qian Yang et al.: Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. undefined, 2018.
- [RGC15] Ribeiro, M.; Grolinger, K.; Capretz, M. A.: MLaaS: Machine Learning as a Service: 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA). IEEE, pp. 896–902, 2015 - 2015.

- [Ru15] Russegger, S. et al.: Big Data und Data-driven Business für KMU. Digital networked Data, 2015.
- [SB14] Shalev-Shwartz, S.; Ben-David, S.: Understanding machine learning. From theory to algorithms. Cambridge University Press, Cambridge, New York, NY, Port Melbourne, Delhi, Singapore, 2014.
- [Sc15] Schryen, G.: Writing Qualitative IS Literature Reviews—Guidelines for Synthesis, Interpretation, and Guidance of Research. Communications of the Association for Information Systems 1/37, 2015.
- [Sc16] Schmarzo, B.: Big data MBA. Driving business strategies with data science. Wiley, Indianapolis, Indiana, 2016.
- [Sc18] Schmarzo, B.: Data Science “Paint by the Numbers” with the Hypothesis Development Canvas. <https://www.linkedin.com/pulse/data-science-paint-numbers-hypothesis-development-canvas-schmarzo/>, accessed 27 Apr 2021.
- [Sc20] Schäfer, F. et al.: Smart Use Case Picking with DUCAR: A Hands-On Approach for a Successful Integration of Machine Learning in Production Processes. Procedia Manufacturing 51, pp. 1311–1318, 2020.
- [SG89] Star, S. L.; Griesemer, J.: Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. Social Studies of Science 3/19, pp. 387–420, 1989.
- [SN17] Sammon, D.; Nagle, T.: The Data Value Map: A framework for developing shared understanding on data initiatives. ECIS 2017: 25th European Conference on Information Systems, pp. 1439–1452, 2017.
- [VD18] VDMA Software und Digitalisierung Ed.: Quick Guide Machine Learning im Maschinen- und Anlagenbau, Frankfurt, 2018.
- [VHR20] Vidgen, R.; Hindle, G.; Randolph, I.: Exploring the ethical implications of business analytics with a business ethics canvas. European Journal of Operational Research 3/281, pp. 491–501, 2020.
- [VLF15] Vossen, G.; Lechtenböcker, J.; Fekete, D.: Big Data in kleinen und mittleren Unternehmen: Eine empirische Bestandsaufnahme (Nr. 135). Arbeitsberichte des Instituts für Wirtschaftsinformatik, Westfälische Wilhelms-Universität Münster, 2015.
- [We19] Weber, C. et al.: A New Process Model for the Comprehensive Management of Machine Learning Models. Proceedings of the 21st International Conference on Enterprise Information Systems ICEIS 2019, pp. 415–422, 2019.
- [WG18] Wolff, D.; Göbel, R.: Digitalisierung. Wie Die Digitalisierung Unsere Lebens- und Arbeitswelt Verändert. Segen oder Fluch. Springer, Berlin, Heidelberg, 2018.
- [WG19] WGP Wissenschaftliche Gesellschaft für Produktionstechnik e.V.: WGP-Standpunkt KI in der Produktion. Künstliche Intelligenz erschliessen für Unternehmen, accessed 7 May 2021.
- [WH00] Wirth, R.; Hipp, J.: CRISP-DM: Towards a standard process model for data mining. Proceedings of the Fourth International Conference on the Practical

Application of Knowledge Discovery and Data Mining, 2000.

- [WW02] Webster, J.; Watson, R. T.: Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly* 2/26, pp. 13–23, 2002.
- [Za19] Zawadzki, J.: Introducing the AI Project Canvas - Towards Data Science. <https://towardsdatascience.com/introducing-the-ai-project-canvas-e88e29eb7024>, accessed 15 Apr 2021.
- [Zh19] Zhou, Z. et al.: ML-Process Canvas: A Design Tool to Support the UX Design of Machine Learning-Empowered Products. *CHI'19 Extended Abstracts*, 2019, Glasgow, Scotland UK, 2019.
- [Zh20] Zhou, Z. et al.: ML Lifecycle Canvas: Designing Machine Learning-Empowered UX with Material Lifecycle Thinking. *Human-Computer Interaction* 5-6/35, pp. 362–386, 2020.