

Machine Learning with Requirements: A Manifesto

Eleonora Giunchiglia^{a,*}, Fergus Imrie^b, Mihaela van der Schaar^{c,e} and Thomas Lukasiewicz^{a,d}

^a *Institute of Logic and Computation, TU Wien, Austria*

E-mails: eleonora.giunchiglia@tuwien.ac.at, thomas.lukasiewicz@tuwien.ac.at

^b *Department of Electrical and Computer Engineering, University of California, Los Angeles, USA*

E-mail: imrie@ucla.edu

^c *DAMTP, University of Cambridge, United Kingdom*

E-mail: mv472@cam.ac.uk

^d *Department of Computer Science, University of Oxford, United Kingdom*

^e *Alan Turing Institute, United Kingdom*

Abstract. In the recent years, machine learning has made great advancements that have been at the root of many breakthroughs in different application domains. However, it is still an open issue how to make them applicable to high-stakes or safety-critical application domains, as they can often be brittle and unreliable. In this paper, we argue that requirements definition and satisfaction can go a long way to make machine learning models even more fitting to the real world, especially in critical domains. To this end, we present two problems in which (i) requirements arise naturally, (ii) machine learning models are or can be fruitfully deployed, and (iii) neglecting the requirements can have dramatic consequences. We show how the requirements specification can be fruitfully integrated into the standard machine learning development pipeline, proposing a novel pyramid development process in which requirements definition may impact all the subsequent phases in the pipeline, and vice versa.

Keywords: Machine Learning, Requirements, Machine Learning Operations, Software Engineering

1. Introduction

In the recent years, machine learning has made great advancements that have been at the root of many breakthroughs in different application domains. For example, AlphaFold [1] is a deep learning model that solved the “protein folding problem”, a grand challenge in the field of biology for more than half a century, while Halicin [2] is the first antibiotic discovered using machine learning, which could help in the battle against bacterial resistance. Results like those above, though ground-breaking, overshadow the dangers that come with the careless use of machine learning models in critical applications. Indeed, even though such models report an astonishingly high performance in terms of accuracy (or alternatively chosen metric), they do not give any guarantee that the model will not have any unintended behaviour when used in practice. Indeed, as pointed out by Rudin in [3], there have already been cases of people erroneously denied parole [4], deep learning models monitoring the air quality in the US that claimed that highly polluted air could be safely breathed [5] and generally poor use of limited resources in medicine, criminal justice, finance and in other domains [6]. Such problems in the behavior of the model are often rooted in the quality of the dataset used for training the model. This is, for example, what happened in the early stages of the Covid-19

*Corresponding author. E-mail: eleonora.giunchiglia@tuwien.ac.at.

1 pandemic, where (as reported in [7]) a common chest scan dataset from [8] made of pediatric scans, was used as
2 control group against Covid-19 positive scans: as a result, instead of Covid-19 detectors, adult/child classifiers were
3 built. Independently from the reasons, such unintended outcomes (i) are particularly dangerous in safety-critical
4 applications, where even a single unforeseen mistake can lead to dramatic consequences, and (ii) undermine human
5 confidence in the models themselves, thus slowing down their adoption.

6 In this paper, we argue that requirements specification and verification can go a long way to make machine learn-
7 ing models even more fitting to the real world, by reducing the risk of getting potentially dangerous unintended
8 behaviors. We start from the simple observation that unintended behaviors correspond to the model violating some
9 requirements which may be known even before the data collection and model development start. To support this
10 claim, we consider two examples, one in healthcare, and the other one in autonomous driving, in which (i) some
11 requirements are known in advance, (ii) machine learning models have a very positive performance in terms of
12 accuracy (or other selected metric), and (iii) despite the positive performance, the outcome of standardly developed
13 machine learning models often violates the spelled requirements with possible dramatic consequences given the crit-
14 icality of the application domains. Though we consider just two examples, we believe that analogous considerations
15 apply to most application domains given (i) the body of knowledge developed along the years in any application
16 domain, which can be translated in corresponding requirements, (ii) the impressive results in performance obtained
17 by machine learning models, which will continue to push for their adoption despite the possible unintended behav-
18 iors, and (iii) the fact that it is not surprising that the resulting model will violate the requirements if they are not
19 taken into account during data collection and model development. Refining the last observation, it is not surprising
20 that a machine learning model has an unintended behavior if, before its deployment, the data and the model itself
21 are not somehow verified against the requirement capturing the intended behavior. We therefore claim that the re-
22 quirements' definition should precede and involve the entire machine learning model development cycle, including
23 the dataset construction. In particular, we propose a novel *pyramid* machine learning development process, which
24 (i) highlights how the requirement definition can help improving every single step of the standard machine learn-
25 ing model development process, and (ii) takes into account that some requirement adjustment might be necessary
26 because of difficulties and/or discoveries that emerge during the development process of machine learning models,
27 especially given their data-driven nature. The above proposal can be seen as a call to adapt and adopt the general
28 methodologies used in software engineering for generic system development, where it is well known (i) that the
29 later the requirements are taken into account the higher the cost is to repair the system, and also (ii) that the require-
30 ment definition is an iterative process in which requirements affect the system development and vice versa, given
31 that during the system development some requirements may be newly discovered and/or adjusted if not canceled
32 (see, e.g., [9–11]). We will ground the above facts in the specific machine learning model development pipeline,
33 highlighting the pros of adopting requirements in the different phases.

34 Our proposal represents a significant shift from the traditional “*performance-driven*” machine learning develop-
35 ment pipeline, which solely concentrates on how to get better performance, measured in terms of accuracy or alter-
36 natively chosen metric. Indeed, we are proposing a “*requirements-driven*” machine learning development pipeline
37 in which performance is just one of the many requirements that the model has to satisfy. Our proposal is thus in
38 line with all the mostly recent works which do not only focus on performance, but also check that the model satis-
39 fies other properties like fairness, robustness, explainability, sustainability, and safety. This obviously might make
40 the development process more complex, because different requirements might be contradictory with one another:
41 a problem that is well known in software engineering [9]. For example, at a high level of abstraction, complex
42 models will likely have a better performance at the possible price of being less explainable and/or verifiable and/or
43 sustainable. Analogously, it might be the case that satisfying certain fairness properties (e.g., equalized odds or
44 demographic parity) is preferable to a higher performance. Independently from the properties that we wish for the
45 model, it is clear that (i) the sooner the requirements are made explicit the better, and (ii) taking into account
46 the requirements in the dataset definition and model construction will likely lead to models that will be easier to
47 verify with respect to the stated requirements. To this end, a good example is given by the recent research in the
48 neuro-symbolic field, where researchers have started developing models that are guaranteed-by-design to be com-
49 pliant with a set of given requirements expressed as logical constraints (see, e.g., [12–14]): in these works, it has
50 been shown how the requirements can be incorporated in the model, which automatically satisfies them, and have a
51 positive impact also on performance.



Fig. 1. Visualization of the standard machine learning pipeline.

Our proposal enjoys the same spirit of works such as [15–17], which advocate for a more structured approach to machine learning model development. In particular, Gebru et al. in [15] propose documentation guidelines for new datasets, Mitchell et al. in [16] advocate for standardized reporting of models, including training data, performance measures, and limitations, while Seedat et al. [17] propose an actionable checklist-style framework to elicit data-centric considerations at different stages of the development pipeline.

The remaining of the paper is structured as follows. First, in Section 2, we consider examples from the healthcare and autonomous driving application domains in order to show that in these domains (i) requirements arise naturally, (ii) machine learning models are or can be fruitfully deployed, and (iii) neglecting the requirements can have dramatic consequences. Secondly, in Section 3, we show how the requirements definition can be fruitfully integrated into the standard machine learning development pipeline, impacting all the phases in the pipeline. Finally, we have the conclusions and possible plans for the road ahead in Section 4.

2. The Need for Requirements

The standard machine learning pipeline is mostly sequential by nature and its main steps are:

1. **Data curation**, in which the dataset for training the model is created, encompassing the following phases: (i) data collection, (ii) data pre-processing, (iii) data augmentation, and (iv) data quality evaluation;
2. **Model Creation**, in which the model is designed and built;
3. **Model Training**, in which the created model is trained and in which the right hyperparameters are chosen (often through testing over a validation dataset);
4. **Model Testing**, in which the performance of the model are assessed (ideally using different metrics and test sets).
5. **Model Deployment**, in which the model is actually deployed in the real world.

A visual representation of the standard sequential pipeline is given in Figure 1. Even though we do not show them in the figure, in practice loop-backs might be necessary in order to modify the dataset or, more often, the model. Still, all such possible changes happen mostly between the model creation step and the model training step, as they are usually driven by the desire to get a better performance, while the dataset is often considered a static element of the pipeline. This is somehow not surprising given that often the unique requirement considered for machine learning models is to have a good performance, according to some selected metrics. However, as highlighted in the introduction, achieving a high score is often not enough to ensure that the model will exhibit the right behaviour when used in practice. Other requirements may be put in place in order to make the model applicable or even certified to be applicable in the given application domain.

We now focus on two application domains, healthcare and autonomous driving, and we show how requirements arise naturally in both applications domains, how applying machine learning techniques has already brought and can bring tremendous advantages to the fields, and how deploying machine learning models in these fields without explicitly taking into account the requirements can lead to unexpected behaviors corresponding to violations of the requirements. Though we consider just two application domains, we believe that analogous considerations and results hold virtually in any application domain, given (i) the body of knowledge developed along the years in any application domain, which can be translated in corresponding requirements, (ii) the impressive results in performance obtained by machine learning models which will continue to push for their adoption despite the possible unintended behaviors, and (iii) the impossibility to certify/rule out the absence of undesired behaviors without explicitly spelling out the corresponding requirements and the testing/verification of the model against the requirements themselves.

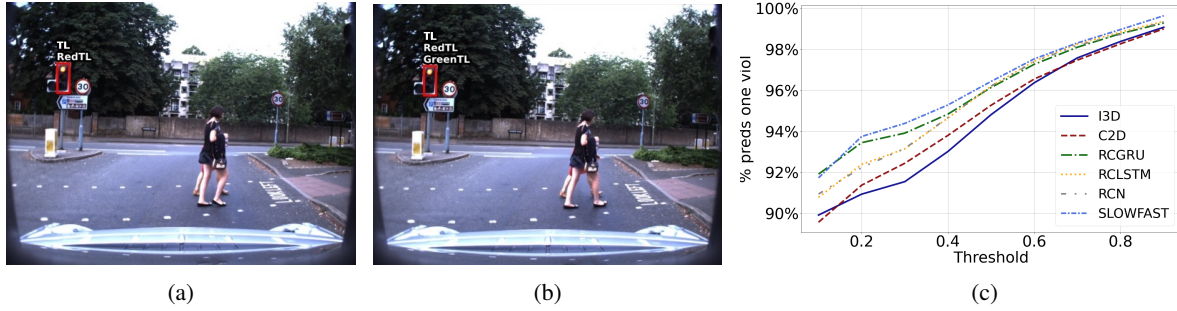


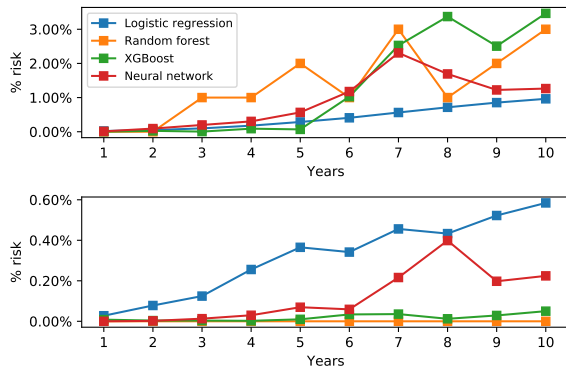
Fig. 2. (a) and (b) show the predictions made by the I3D model (with $\theta = 0.5$) for the same traffic light and just one frame apart. (c) shows the number of predictions that violate at least one requirements when varying θ ; (c) is from [14].

2.1. Autonomous Driving

In the recent years, the developments in machine learning, and in particular in computer vision, have fuelled the hopes of autonomous vehicles being finally in reach. However, every so often, such hopes are crushed by car crashes that, in some cases, have even injured or even killed people. For example, in March 2018, an autonomous vehicle developed and tested on public roads by Uber’s Advanced Technologies Group fatally injured a pedestrian who was pushing their bicycle across the street outside of a designated crossing area [18]. The most striking characteristic of this accident is the fact that the car did not make any attempt to break and/or to avoid the pedestrian [19]. Unfortunately, this is not an isolated incident, as nearly 400 car crashes involving autonomous vehicles have been reported in the United States over a period of only ten months in 2022 [20]. As reported in [20], these vehicles in order to take their decisions rely, among others, on computer vision models, which if trained to simply maximise their performance (following the standard machine learning development pipeline) might fail to abide to even the simplest requirements.

To exemplify the problem described above, we consider the ROAD-R dataset [14], which consists of (i) 22 relatively long (~ 8 minutes each) videos annotated with road events, i.e., a sequence of frame-wise bounding boxes linked in time, each labelled with the agent in the bounding box, together with its action(s) and location(s), and (ii) 243 requirements expressed in propositional logic stating what is an admissible road event.¹ The possible labels associated to each bounding box are 41, and the requirements state which combinations of labels a model can output. Hence, for instance, a requirement states that a traffic light cannot be red and green at the same, while another states that a traffic light cannot move. Given ROAD-R, six state-of-the-art temporal feature learning architectures (I3D [21], C2D [22], RCGRU [23], RCLSTM [23], RCN [24] and SlowFast [25]) as part of a 3D-RetinaNet model [26] (with a 2D-ConvNet backbone made of Resnet50 [27]) for event detection, have been trained. These models take as input a set of consecutive frames, and for each frame they output (i) a set of bounding boxes, and (ii) a set of labels for each bounding box. Such labels are decided in the standard way: for each of the 41 labels the model outputs a number $o \in [0, 1]$, and if $o > \theta$ then the label is returned, otherwise it is not. θ is a user-defined threshold which is often picked equal to 0.5. An example of prediction is given in Figure 2a, where the prediction for the depicted bounding box is {Traffic Light, Red Traffic Light}. The surprising finding of the work [14] is that, as shown in Figure 2c, no matter the chosen threshold θ , at least 89% of the predictions violate at least one requirement. Even more, some of the predictions may violate requirements corresponding to common knowledge possessed by humans, making the prediction difficult to interpret and manage. For example, if we consider the prediction in Figure 2b, done by the same model that made the prediction in Figure 2a, for the same traffic light just one frame later, then we see that not only the prediction is wrong, but it also violates the common knowledge (in this case corresponding also to a formally stated requirement) that the traffic light cannot be red and green at the same time. Such prediction, if not further elaborated by the system controlling the vehicle could indeed have dramatic consequences. For this reason,

¹Dataset available at: <https://github.com/gurkirt/road-dataset>. Requirements available at: <https://github.com/EGiunchiglia/ROAD-R>.



(a) Example predictions at different time horizons.


(b) Proportion of patients with ≥ 1 violation.

Fig. 3. Figure 3a shows two examples of the predictions made by different classification models for the same patient at different time horizons. Figure 3b shows the proportion of samples for which the predictions violate the requirement for risk to be increasing for longer time horizons.

it is clear that any such prediction will have to be specifically handled by the system controlling the vehicle, which for sure will incorporate and satisfy the violated requirements. However, if the requirements are incorporated into and surely verified by the entire system (as it must be the case in this setting), it is far from being clear why they are standardly neglected in the process of building the machine learning component of the system. Indeed, none of the six evaluated systems is able to handle requirements on their inputs-outputs. Even more, as reported in [14], the dataset used for training the model violates some of the constraints (like the fact that it is not possible for a vehicle to be both incoming and outgoing), which surely will have to be incorporated into any real application of the models.²

2.2. Healthcare

For our second example, we discuss healthcare. There is great hope that developments in machine learning will revolutionize medicine and transform clinical practice [28]. The range of applications in medicine is vast, from computer vision systems analyzing medical images in radiology and longitudinal monitoring of patient trajectories throughout a hospital stay, to genomic screening of future disease risk and much more. However, as with many other high-stakes and safety-critical applications of machine learning, there are a number of requirements for machine learning systems in healthcare. These include ethical considerations, such as fairness and bias; practical considerations, such as controlling false positive rate to prevent alarm fatigue [29] and ensuring appropriate resource allocation; and logical considerations, such as the example discussed below, among others.

As a concrete example, we consider clinical risk scores. Clinical risk scores estimate the likelihood of a specific outcome occurring after a certain period of time, such as a patient developing a particular disease or condition in the next ten years or experiencing an adverse event following a medical procedure. By definition, the chance of an outcome occurring within some time horizon t must increase with time, thus risk scores must be monotonically increasing functions of time, i.e., the predicted n -year risk is greater than their m -year risk, for all $n > m$. However, there are many studies that use classification models to predict risk after a specific time horizon, for example, 5-year risk. Naturally, we can generalize this by predicting risk at multiple time horizons rather than just one. Often this is very useful clinically, since patient trajectories can vary greatly and patients with equivalent long-term risks might have divergent short-term risks. A straightforward way to achieve this is to train multiple classification models to predict the risk at each time point of interest. However, there is no guarantee that the predicted risk will increase for longer time horizons.

²While it might be argued that it is normal to have errors in the dataset because of noise, it seems odd to train a model with known-to-be-wrong-data and then correct the model outputs when they respect the known-to-be-wrong-data.

To illustrate the consequence of not incorporating the simple requirement described above into the machine learning systems, we consider the problem of predicting the risk of developing diabetes. We construct a cohort of patients from UK Biobank [30], a large-scale observational study with around 500,000 participants from 22 assessment centers across England, Wales, and Scotland enrolled between 2006 and 2010. We extracted a cohort of participants who were 40 years of age or older at enrollment with no diagnosis or history of diabetes at baseline. We considered the 18 features employed by QDiabetes [31]. We performed data imputation using HyperImpute [32] and trained classification algorithms using AutoPrognosis [33] to predict n -year risk of developing diabetes for $n \in \{1, \dots, 10\}$. We averaged results over five random initializations. Two examples are provided in Figure 3a. For the first patient (Figure 3a, top), only the logistic regression models met the requirement that the risk of developing diabetes should be monotonically increasing over longer time horizons, with the random forest, XGBoost, and neural network models having at least one time horizon with lower predicted risk than the previous. A similar situation can be seen for the second patient (Figure 3a, bottom), where only the random forest model satisfies the requirement. Considering the entire dataset and all ten time horizons, around 20% of patients have predictions that violate the requirement for risk to monotonically increase at least once for the logistic regression and random forest models, while for over 99% of patients XGBoost and neural networks issued predictions that violated the requirement. Note that the number of violations made by the models was not necessarily correlated with performance, with the lowest-performing model as measured by area under the receiver operating curve at the 5-year horizon exhibiting the fewest violations (random forests). Such predictions, at best, are simply inaccurate but, at worst, can erode trust in machine learning systems and have more serious consequences depending on the actions taken. There are several solutions to satisfy this particular requirement, such as specialized functional forms [34, 35], multi-task learning with specialized loss functions [36], or bespoke architectures. While the direct real-world impact of such predictions is likely limited for this particular example, the issue and need for requirements is clear.

3. The Pyramid Model

In the previous section, we have shown the standard performance-driven machine learning development pipeline, in which the requirements are traditionally neglected, and, with the means of two examples, we have highlighted how such negligence can have dire consequences. This automatically calls for the inclusion of the requirements in the development process, which are at the core of any development pipeline proposed in the field of software engineering. We thus propose to adapt the general methodologies used in software engineering for generic system development, where it is well known that the later the requirements are taken into account the higher the cost is to repair the system. In order to adapt such methodologies, we need to take into account that, differently from standard software, machine learning models are learnt from data, and thus we have lower control over the behaviour of the model, which entails not only that at design time it is possible to predict the future behaviour of the model, but also, at testing time, if a bug is found, then the steps necessary to fix it are entirely different. A development pipeline for a machine learning model thus needs to take into account that:

1. data are central to the process, as the quality of the model heavily depends on the quality of the collected data,
2. the model is learnt from data, and thus it is not possible to fully know a priori its behaviour, and
3. if the model does not behave as expected, it is advisable to not only check the model but also check the data.

Further, it is unrealistic to expect that all requirements are always known a priori and that they remain immutable throughout the development process. This is already known in the general software engineering field, and it has been already specifically tested to be the case also in our field. Indeed, as highlighted by the survey conducted in [37], some of the biggest problems that machine learning practitioners face on a daily basis include the accessibility of data (which might not be known in the requirements definition phase) and the unrealistic expectations of the stakeholders of the system (which might reveal themselves during the testing phase of the model). Thus, we also need to take into account that the requirements might change at each stage of the development process.

Given the above, we propose the *pyramid machine learning development model* in Figure 4, in which the requirement definition phase is (as expected) put at the beginning of the development process and in which the pyramid shape allows to better highlight the close relationship that must exist between the requirement definition phase and




Fig. 4. Visualization of the pyramid model. The full arrows stand for the standard procedural processes, while the dotted arrows show that the requirements impact every stage of the process.

all the others in the pipeline. Consider the grey backward lines. These represent the fact that if something unexpected happens at any point of the pipeline (e.g., the model does have some unintended behavior during the testing phase), then it is necessary to go back and first check if the defined requirements were suitable or need to be updated, then if the collected data are representative of the phenomenon that we want to capture, and so on. Obviously, the later in the pipeline we realise that there is a bug, the more costly it is, and this is captured in Figure 4 by the length of the path to cover as we advance in the pipeline. Consider now the dotted arrows, which illustrate the fact that the requirements can help in shaping each of the development steps, and in particular:

1. **Data Curation:** the requirements over the data obviously define the data collection phase, as they state which properties the dataset should have. This can positively impact this phase as it forces the stakeholders of the system to carefully reflect on which aspects of the total population the dataset should capture and which should be ignored. Furthermore, they can help in reducing the mistakes done in the annotation phase (as any annotation that is not compliant with the requirements has to be corrected) and they can even fasten it (as the annotators can be shown only those options that are compliant with the requirements, or they can be automatically corrected in real-time).
2. **Model Creation:** the requirements over the model define how the model should behave, and thus they should be taken into account at creation time. Following the same principles used for the data curation step, having requirements can positively impact this phase, because it forces the stakeholders of the system to reflect upon its expected behaviour and which outcomes must be avoided at all times. Such requirements can then be mapped in the model itself. As an example, given a set of formal requirements, some works [12, 13, 38] in the neuro-symbolic field were able to map the requirements directly to the topology of neural networks and guarantee that the constraints are always satisfied, while exploiting the background knowledge expressed by the requirements to even get better performance.
3. **Model Training:** the requirements over the model can again be used to define the objectives of our training. This is what it is done for example in many works in the fairness field, where the model is trained with two objectives: one to maximise performance and one to maximise the desired fairness definition (see, e.g., [39]). Furthermore, if we map the requirements to the loss function, we can even use them to alleviate the need for labelled data in semi-supervised settings (see, e.g., [40–42]). Indeed, given an unlabelled data point, the model can be taught to simply return an output which is compliant with the requirements.

4. **Model Testing:** as done in standard software engineering, the requirements should guide the testing phase. Each requirement should be checked to hold (eventually using formal verification techniques; see, e.g., [43, 44]) and in case it does not, then either the requirement has to be further analyzed and eventually modified or some procedure has to be put in place in order to signal that the requirement has to be properly handled outside of the machine learning model.
5. **Model Deployment:** requirements can shape also the deployment phase, especially if the stakeholders want the model to retain certain properties even in presence of data shift and/or drift.

4. Summary and Outlook

In this paper, we have shown that even though applying machine learning to a given problem can bring great benefits, applying it without specifying and taking into account the requirements of the problem can lead to unexpected behaviors and, possibly, dramatic consequences. We have thus proposed a new machine learning development pipeline in which the requirements definition phase is explicitly incorporated at the beginning of the process, causing a deep connection between the requirements definition and all the other phases in the development process, in which (as standard in software engineering) the former may affect the latter and vice versa. From this perspective, this paper can be seen as a call to adopt the general methodologies used in software engineering for generic system development to the specific field of machine learning, highlighting the risks of not doing so. We believe that in many cases the benefits of explicitly defining the requirements and taking them into account in the other phases outweigh their cost. It is also clear that the requirement definition and exploitation is unavoidable when machine learning models are used as stand-alone systems in domains with strict requirements to satisfy (e.g., in safety critical domains) or when used as part of larger systems with a given set of stringent requirements (e.g., memory usage; see, e.g., [45]). This view is supported by the recent trends in machine learning in which performance is just one of the requirements to satisfy. In particular, the works done in the fields of interpretability (see, e.g., [46, 47]), fairness (see, e.g., [48, 49]) and robustness (see, e.g., [50, 51]) all privilege certain characteristics of the model over performance, and all these works roots their motivation in the necessity to take into account requirements emerging from the respective application domain. However, how to define and fully exploit the domain knowledge expressed by requirements in the development process is still largely an open question, as many technical challenges need to be overcome. To this end, the neuro-symbolic field can be seen as a bridge between standard machine learning and software engineering, as researchers in the field have already started developing techniques to create models that are built not only by the data, but that take explicitly as input also other type of knowledge, like formally stated requirements and use them either to design the model (see, e.g., [12, 13]), or to incorporate the requirements in the loss function (see, e.g., [41, 42]).

Despite the huge amount of work that can be categorized under the umbrella of “requirement-driven” machine learning, this is the first paper which explicitly advocates for a general model development methodology in which the requirement definition is a first class citizen like the other phases in the pipeline. We believe that this corresponds to a more structured approach to machine learning model development, which has already been advocated in [15–17].

References

- [1] A.W. Senior et al., Improved protein structure prediction using potentials from deep learning, *Nature* **577**(7792) (2020). ISBN 1476-4687.
- [2] J.M. Stokes et al., A Deep Learning Approach to Antibiotic Discovery, *Cell* **180**(4) (2020).
- [3] C. Rudin, Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nature Machine Intelligence* **1**(5) (2019).
- [4] R. Wexler, When a computer program keeps you in jail: How computers are harming criminal justice., *The New York Times* (2017). <https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice.html>.
- [5] M. McGough, How bad is Sacramento’s air, exactly? Google results appear at odds with reality, some say, *Sacramento Bee* (2018). <https://www.sacbee.com/news/california/fires/article216227775.html>.
- [6] K.R. Varshney and H. Alemzadeh, On the Safety of Machine Learning: Cyber-Physical Systems, Decision Sciences, and Data Products, *Big data* **5**(3) (2016).

- [7] M. Roberts, D. Driggs, M. Thorpe, J. Gilbey, M. Yeung, S. Ursprung, A.I. Aviles-Rivero, C. Etmann, C. McCague, L. Beer, J.R. Weir-McCall, Z. Teng, E. Gkrania-Klotsas, A. Ruggiero, A. Korhonen, E. Jefferson, E. Ako, G. Langs, G. Gozaliasl, G. Yang, H. Prosch, J. Preller, J. Stanczuk, J. Tang, J. Hofmanninger, J. Babar, L.E. Sánchez, M. Thillai, P.M. Gonzalez, P. Teare, X. Zhu, M. Patel, C. Cafolla, H. Azadbakht, J. Jacob, J. Lowe, K. Zhang, K. Bradley, M. Wassin, M. Holzer, K. Ji, M.D. Ortet, T. Ai, N. Walton, P. Lio, S. Stranks, T. Shadbahr, W. Lin, Y. Zha, Z. Niu, J.H.F. Rudd, E. Sala, C.-B. Schönlieb and AIX-COVNET, Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans, *Nature Machine Intelligence* **3**(3) (2021), 199–217.
- [8] D. Kermany, M. Goldbaum, W. Cai, C. Valentim, H.-Y. Liang, S. Baxter, A. McKeown, G. Yang, X. Wu, F. Yan, J. Dong, M. Prasadha, J. Pei, M. Ting, J. Zhu, C. Li, S. Hewett, J. Dong, I. Ziyar and K. Zhang, Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning, *Cell* **172** (2018), 1122–1131.e9.
- [9] I. Sommerville and P. Sawyer, *Requirements Engineering: A Good Practice Guide*, Wiley, 1997.
- [10] I. Sommerville, *Software engineering, 10th Edition*, International computer science series, Addison-Wesley, 2015.
- [11] B. Nuseibeh and S.M. Easterbrook, Requirements engineering: a roadmap, in: *Proc. of ICSE*, A. Finkelstein, ed., ACM, 2000, pp. 35–46.
- [12] E. Giunchiglia and T. Lukasiewicz, Multi-Label Classification Neural Networks with Hard Logical Constraints, *JAIR* **72** (2021).
- [13] N. Hoernle, R. Karampatsis, V. Belle and K. Gal, MultiplexNet: Towards Fully Satisfied Logical Constraints in Neural Networks, in: *Proc. of AAAI*, 2022.
- [14] E. Giunchiglia, S. Khan, F. Cuzzolin and T. Lukasiewicz, ROAD-R: The Autonomous Driving Dataset for Learning with Requirements, *Machine Learning Journal* (2023).
- [15] T. Gebru, J. Morgenstern, B. Vecchione, J.W. Vaughan, H.M. Wallach, H.D. III and K. Crawford, Datasheets for Datasets, *CoRR abs/1803.09010* (2018). <http://arxiv.org/abs/1803.09010>.
- [16] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I.D. Raji and T. Gebru, Model Cards for Model Reporting, in: *Proc. of FAT**, danah boyd and J.H. Morgenstern, eds, ACM, 2019, pp. 220–229.
- [17] N. Seedat, F. Imrie and M. van der Schaar, DC-Check: A Data-Centric AI checklist to guide the development of reliable machine learning systems, *CoRR abs/2211.05764* (2022). doi:10.48550/arXiv.2211.05764.
- [18] NTSB, Preliminary report: Highway HWY18MH010., 2018, Washington, DC: National Transportation Safety Board.
- [19] C. Macrae, Learning from the Failure of Autonomous and Intelligent Systems: Accidents, Safety, and Sociotechnical Sources of Risk, *Risk analysis : an official publication of the Society for Risk Analysis* **42** (2022).
- [20] N.E. Boudette, C. Metz and J. Ewing, Tesla Autopilot and Other Driver-Assist Systems Linked to Hundreds of Crashes, *The New York Times* (2022). <https://www.nytimes.com/2022/06/15/business/self-driving-car-nhtsa-crash-data.html>.
- [21] J. Carreira and A. Zisserman, Quo vadis, action recognition? A new model and the kinetics dataset, in: *Proc. of CVPR*, 2017.
- [22] X. Wang, R. Girshick, A. Gupta and K. He, Non-local neural networks, in: *Proc. of CVPR*, 2018.
- [23] Y. Hua, Z. Zhao, Z. Liu, X. Chen, R. Li and H. Zhang, Traffic prediction based on random connectivity in deep learning with long short-term memory, in: *Proc. of VTC-Fall*, 2018.
- [24] G. Singh and F. Cuzzolin, Recurrent convolutions for causal 3D-CNNs, in: *Proc. of ICCV Workshops*, 2019.
- [25] C. Feichtenhofer, H. Fan, J. Malik and K. He, SlowFast networks for video recognition, in: *Proc. of ICCV*, 2019.
- [26] G. Singh, S. Akrigg, M.D. Maio, V. Fontana, R.J. Alitappeh, S. Saha, K.J. Saravi, F. Yousefi, J. Culley, T. Nicholson, J. Omokeowa, S. Khan, S. Grazioso, A. Bradley, G.D. Gironimo and F. Cuzzolin, ROAD: The ROad event Awareness Dataset for Autonomous Driving, *IEEE TPAMI* (2022).
- [27] K. He, X. Zhang, S. Ren and J. Sun, Deep residual learning for image recognition, in: *Proc. of CVPR*, 2016.
- [28] E.J. Topol, High-performance medicine: The convergence of human and artificial intelligence, *Nature Medicine* **25**(1) (2019), 44–56. doi:10.1038/s41591-018-0300-7.
- [29] S. Sendelbach and M. Funk, Alarm fatigue: A patient safety concern, *AACN Advanced Critical Care* **24**(4) (2013), 378–386.
- [30] C. Sudlow, J. Gallacher, N. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. Landray, B. Liu, P. Matthews, G. Ong, J. Pell, A. Silman, A. Young, T. Sprosen, T. Peakman and R. Collins, UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age, *PLOS Medicine* **12**(3) (2015), 1–10. doi:10.1371/journal.pmed.1001779.
- [31] J. Hippisley-Cox and C. Coupland, Development and validation of QDiabetes-2018 risk prediction algorithm to estimate future risk of type 2 diabetes: cohort study, *BMJ* **359** (2017). doi:10.1136/bmj.j5019.
- [32] D. Jarrett, B.C. Cebere, T. Liu, A. Curth and M. van der Schaar, HyperImpute: Generalized Iterative Imputation with Automatic Model Selection, in *Proceedings of Machine Learning Research*, Vol. 162, PMLR, 2022, pp. 9916–9937.
- [33] F. Imrie, B. Cebere, E.F. McKinney and M. van der Schaar, AutoPrognosis 2.0: Democratizing Diagnostic and Prognostic Modeling in Healthcare with Automated Machine Learning, *arXiv preprint arXiv:2210.12090* (2022).
- [34] D.R. Cox, Regression Models and Life-Tables, *Journal of the Royal Statistical Society: Series B (Methodological)* **34**(2) (1972), 187–202. doi:<https://doi.org/10.1111/j.2517-6161.1972.tb00899.x>.
- [35] C. Lee, W. Zame, J. Yoon and M. Van Der Schaar, DeepHit: A deep learning approach to survival analysis with competing risks, in: *Proc. of AAAI*, Vol. 32, 2018.
- [36] Y. Li, J. Wang, J. Ye and C.K. Reddy, A multi-task learning formulation for survival analysis, in: *Proc. of ACM SIGKDD*, 2016, pp. 1715–1724.
- [37] S. Mäkinen, H. Skogström, E. Laaksonen and T. Mikkonen, Who Needs MLOps: What Data Scientists Seek to Accomplish and How Can MLOps Help?, in: *1st IEEE/ACM Workshop on AI Engineering - Software Engineering for AI*, IEEE, 2021, pp. 109–112.
- [38] K. Ahmed, S. Teso, K. Chang, G.V. den Broeck and A. Vergari, Semantic Probabilistic Layers for Neuro-Symbolic Learning, in: *Proc. of NeurIPS*, 2022.

- [39] A. Agarwal, A. Beygelzimer, M. Dudík, J.L. 0001 and H.M. Wallach, A Reductions Approach to Fair Classification, in: *Proc. of ICML*, PMLR, 2018, pp. 60–69.
- [40] J. Xu, Z. Zhang, T. Friedman, Y. Liang and G. Van den Broeck, A Semantic Loss Function for Deep Learning with Symbolic Knowledge, in: *Proc. of ICML*, 2018.
- [41] M. Diligenti, M. Gori, M. Maggini and L. Rigutini, Bridging logic and kernel machines, *Mach. Learn.* **86** (2012).
- [42] S. Badreddine, A. d’Avila Garcez, L. Serafini and M. Spranger, Logic Tensor Networks, *Artif. Intell.* **303** (2022).
- [43] L. Pulina and A. Tacchella, An Abstraction-Refinement Approach to Verification of Artificial Neural Networks, in: *Proc. of CAV*, Lecture Notes in Computer Science, Vol. 6174, Springer, 2010, pp. 243–257.
- [44] A. Lomuscio and L. Maganti, An approach to reachability analysis for feed-forward ReLU neural networks, *CoRR abs/1706.07351* (2017). <http://arxiv.org/abs/1706.07351>.
- [45] M. Rhu, N. Gimelshein, J. Clemons, A. Zulfiqar and S.W. Keckler, vDNN: Virtualized deep neural networks for scalable, memory-efficient neural network design, in: *Proc. of IEEE/ACM MICRO*, IEEE Computer Society, 2016, pp. 18:1–18:13.
- [46] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin and J. Su, This Looks Like That: Deep Learning for Interpretable Image Recognition, in: *Proc. of NeurIPS*, H.M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E.B. Fox and R. Garnett, eds, 2019, pp. 8928–8939.
- [47] F. Doshi-Velez and B. Kim, Towards A Rigorous Science of Interpretable Machine Learning, 2017.
- [48] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman and A. Galstyan, A Survey on Bias and Fairness in Machine Learning, *ACM Comput. Surv.* **54**(6) (2022), 115:1–115:35. doi:10.1145/3457607.
- [49] M. Wick, S. Panda and J.-B. Tristan, Unlocking Fairness: a Trade-off Revisited, in: *Proc. of NeurIPS*, Vol. 32, Curran Associates, Inc., 2019.
- [50] T. Weng, H. Zhang, P. Chen, J. Yi, D. Su, Y. Gao, C. Hsieh and L. Daniel, Evaluating the Robustness of Neural Networks: An Extreme Value Theory Approach, in: *Proc. of ICLR*, OpenReview.net, 2018.
- [51] G. Katz, C.W. Barrett, D.L. Dill, K. Julian and M.J. Kochenderfer, Towards Proving the Adversarial Robustness of Deep Neural Networks, in: *Proceedings First Workshop on Formal Verification of Autonomous Vehicles, FVAV@iFM 2017, Turin, Italy, 19th September 2017*, L. Bulwahn, M. Kamali and S. Linker, eds, EPTCS, Vol. 257, 2017, pp. 19–26. doi:10.4204/EPTCS.257.3.
- [52] P. Dragone, S. Teso and A. Passerini, Neuro-Symbolic Constraint Programming for Structured Prediction, in: *IJCLR-NeSy*, 2021.
- [53] X. Huang, D. Kroening, W. Ruan, J. Sharp, Y. Sun, E. Thamo, M. Wu and X. Yi, A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability, *Comput. Sci. Rev.* **37** (2020), 100270. doi:10.1016/j.cosrev.2020.100270.
- [54] C. Cheng, C. Huang and H. Yasuoka, Quantitative Projection Coverage for Testing ML-enabled Autonomous Systems, in: *Proc. of ATVA*, Lecture Notes in Computer Science, Vol. 11138, Springer, 2018, pp. 126–142.
- [55] A. Odena, C. Olsson, D.G. Andersen and I.J. Goodfellow, TensorFuzz: Debugging Neural Networks with Coverage-Guided Fuzzing, in: *Proc. of ICML*, Proceedings of Machine Learning Research, Vol. 97, PMLR, 2019, pp. 4901–4911.
- [56] W.W. Royce, Managing the Development of Large Software Systems: Concepts and Techniques, in: *Proc. of ICSE*, ACM Press, 1987, pp. 328–339.
- [57] B.W. Boehm, A Spiral Model of Software Development and Enhancement, *Computer* **21**(5) (1988), 61–72. doi:10.1109/2.59.
- [58] G. Symeonidis, E. Nerantzis, A. Kazakis and G.A. Papakostas, MLOps - Definitions, Tools and Challenges, in: *Proc. of IEEE CCWC*, IEEE, 2022, pp. 453–460.
- [59] M.M. John, H.H. Olsson and J. Bosch, Towards MLOps: A Framework and Maturity Model, in: *Proc. of SEAA*, IEEE, 2021, pp. 1–8.
- [60] N. Eén and N. Sörensson, An Extensible SAT-solver, in: *Proc. of SAT*, Lecture Notes in Computer Science, Vol. 2919, Springer, 2003, pp. 502–518.