A conceptual model for leaving the data-centric approach in machine learning

SEBASTIAN SCHER, Know-Center, Austria

BERNHARD GEIGER, Know-Center, Austria

SIMONE KOPEINIK, Know-Center, Austria

ANDREAS TRÜGLER, Know-Center, Austria, Institute of Interactive Systems and Data Science, Graz University of Technology, Austria, and Department of Geography and Regional Science, University of Graz, Austria

DOMINIK KOWALD, Know-Center, Austria and Institute of Interactive Systems and Data Science, Graz University of Technology, Austria

For a long time, machine learning (ML) has been seen as the abstract problem of learning relationships from data independent of the surrounding settings. This has recently been challenged, and methods have been proposed to include external constraints in the machine learning models. These methods usually come from application-specific fields, such as de-biasing algorithms in the field of fairness in ML or physical constraints in the fields of physics and engineering. In this paper, we present and discuss a conceptual high-level model that unifies these approaches in a common language. We hope that this will enable and foster exchange between the different fields and their different methods for including external constraints into ML models, and thus leaving purely data-centric approaches.

1 INTRODUCTION

The essence of supervised Machine Learning (ML) is to learn the relationship in paired data. For a long time, the development of ML models has been seen as an abstract problem that can rely on data alone, and is agnostic of the surrounding setting. However, in many settings, the predictions made by models that solely rely on data are unsatisfactory in practical applications. Here, we on purpose use the generic word "unsatisfactory" instead of more technical terms such as "not accurate", as such terms are usually associated with being computed solely from the data. Cases of predictions being unsatisfactory could for example be unfair predictions —e.g. predictions that are biased towards a specific gender—, or predictions that violate physical laws. In order to make the predictions satisfactory, one can, in addition to the data, introduce external constraint(s). Here we will deal with this idea on a very high conceptual level that is application- and field-agnostic, and thus can be applied to a wide range of settings. This in contrast to more technical discussions that focus on a particular problem set (e.g., including knowledge about physics in a physics setting).

We develop a conceptual model for the purpose of ensuring that an ML model adheres to constraints that are not (or not necessarily) represented by the data from which the ML model is learned. A related approach is "Seldonian ML" [24], which is a way to include probabilistic constraints into ML applications in order to prevent "undesirable behaviour". The approach of Seldonian ML is generic, and can be used for different domains. It can be seen as a mathematical framework that combines requirements for the algorithm specified by the algorithm designer with constraints specified by the user of the algorithm. Our paper employs a view on a much higher, non-mathematical, more conceptual level.

In the first part of the paper, we develop and present the conceptual model. In the second part, we show how it can be applied in two different fields, for both of which there is a large body of research on ensuring that the ML models adheres to certain constraints: physics-informed ML on the one hand, and fair ML on the other hand. While the constraints in these two fields come from very different perspectives —physical laws in the case of physics-informed ML, and ethical constraints in the case of fair ML — we show that with our conceptual model it is possible to describe them in the same language. The goal of this study is not to offer concrete solutions to practitioners, but to introduce a high-level view that can help in conceptually transferring methods between different domains.

2 COMMON CONCEPTUAL MODEL

In this section, we describe the conceptual model that can serve as a common framework for leaving the data-centric approach in ML. We will use the following terms and notations:

- the ML model M
- the dataset D
- the worldview W
- the constraints C

The ML model M is very general here and can be the result of any type of supervised ML algorithm, including combinations of different algorithms. The ML model M in our discussion also subsumes pre- and post-processing steps. The ML model trained with the data D is the raw ML model, which we will denote M_r. If the training algorithm is appropriate, this model will make predictions that adequately represent the relations in D. Whether this implies successful generalization, or whether this requires certain train/validation set splits during the learning phase, is immaterial for the remainder of this work, which is why we omit such discussions. However, the predictions of M_r need not be *satisfactory* in an actual setting, in which D is only part of the whole thing. In addition, the worldview W comes in. Our worldview describes how things ought to be, independent of what the data says. The worldview is thus both required and is essential for leaving the data-centric approach. Critically, W can be in stark conflict with D. W can be broad and generic, such as "A and B have to be treated equally", or "physical law X has to be observed". The formulation of W is not specific to a certain ML model. Instead, for a specific ML model, it must be translated into constraints C. The exact form of C will depend on the application and the algorithm used for training. The constraints could either be hard constraints (such as architectural choices, or constraints in classical optimization), or soft constraints (such as regularization, or additional objectives in classical optimization) – in the latter case one could also refer to them as a "cost". In our discussion, the term constraints C is used in a generic way and covers both types of constraints. These constraints are then incorporated in the ML Model M. This results in the constrained ML Model M_C. M_C should then provide *satisfactory* results with respect to W. The conceptual model is illustrated in fig. 1. Formalizing W, translating it into C, and incorporating C into M requires additional knowledge on top of D, and the exact process will vary vastly between different fields, applications and algorithms. This is, however, the essential step for leaving the data-centric paradigm. In the next sections, this will be illustrated with the two fields of physics-informed ML and fair ML.

3 EXAMPLES OF APPLYING THE CONCEPTUAL MODEL

In this section, we first briefly discuss the domains of physics-informed ML and fair ML, and how to approach them with our conceptual model. Then, in less detail, multiple other domains are discussed.
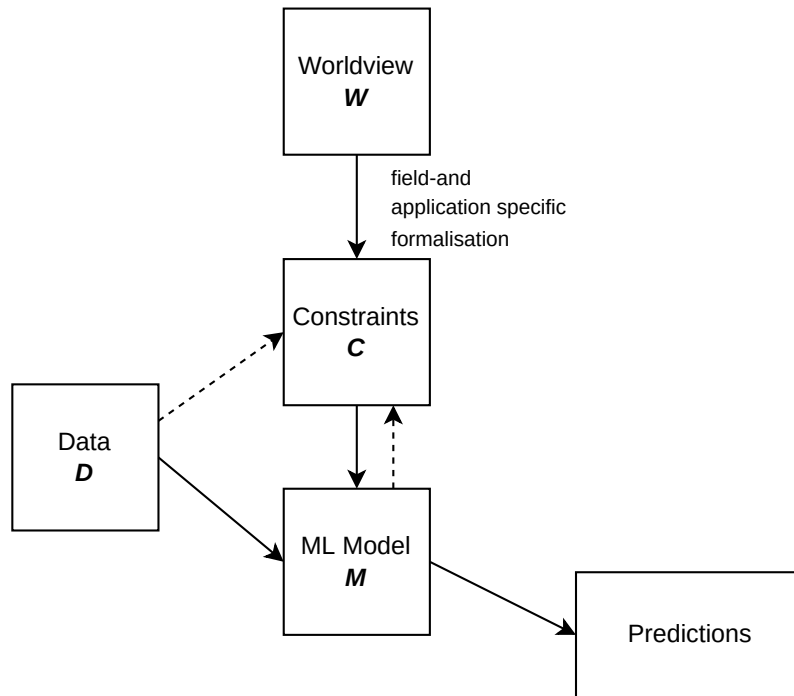
Fig. 1. Illustration of the conceptual model: The worldview W forms the basis for the constraints C, and is independent of the ML Model M and the data D. C is the translation of W for a given application, data and ML model, and thus also depends on D and M (indicated by the dotted arrows). The resulting ML Model and predictions are then influenced both by D and W.

3.1 W given by physics - physics informed machine learning

The goal of physics-informed ML (PIML) is to include information about physical laws into ML models. The most widespread form of PIML are Physics-Informed Neural Networks (PINNs) [19], which incorporate physical laws via sets of ordinary or partial differential equations and which are used to solve initial/boundary value and inverse problems. Here we summarize only the basic concepts of PIML, for a comprehensive overview see [10]. There are two main approaches towards PIML:

(1) Enforcing physical laws (e.g. certain invariances) with specially designed ML approaches, for example with neural network architectures that automatically follow a certain physical law, or kernels that are given by a certain physical task. Examples include a variant of PINNs, in which initial and boundary values are enforced using predefined functions [16], or graph neural networks which are trained such that the kernel resembles a physical operator in mesh-based simulations [20].

(2) Including physical laws in the cost function of M (e.g. [7, 11, 23]). With this, the ML model is penalized in the learning phase if it does not respect the physical law, but it is not actually forced to obey it, and there is no guarantee that it is obeyed. This is a soft approach, in which the physical law(s) are implicitly weighted against other things (e.g. prediction performance). The big upside is that it is generally much easier to include rules/laws in the cost function than enforcing them through specialized architectures/models, and that the same cost function can be used for many different methods.

In high-level language, in problem settings were PIML is a possible solution we are faced with the following situation: we have data D that approximates the real world, but it is not perfect and can have all types of errors (such as measurement bias, sensor drift, noise, etc). These errors may manifest themselves as violations of physical laws in D, such as an energy loss over time, even if the data is from a system where energy is conserved. If we train an ML model M_r with this raw data only, its predictions will have errors and therefore be unsatisfactory. At least part of these errors will originate from the fact that D has errors. With physics-constrained ML, we incorporate a-priori physical knowledge in order to make the predictions more satisfactory. In this context, W is given by laws of physics, e.g., systems of differential equations or conservation laws that describe that a certain quantity (energy, momentum, mass, etc.) cannot change. Based on W we thus tell the algorithm "this is how it should be", because these are the physical laws governing the problem at hand. These form the constraints C, and making use of the various available methods of PIML, can be used to create the desired constrained model M_C. An illustration of incorporating physics laws in our conceptual model is shown in the left panel of Fig. 2.

3.2 W given by ethical constraints - fair machine learning

Another domain where the predictions of the raw model M_r are often unsatisfactory are settings that - either directly or indirectly - involve people. This is the topic of research on algorithmic fairness, and is often alternatively referred to as fair ML, or more colloquially as "bias" problems. In settings with fairness issues, from a high-level viewpoint, the problem is that "how it should be" is harder to define compared to, e.g., constraints by laws of physics. This is because the worldview W, which defines the ethical system that one applies, is not as explicit and quantifiable as a worldview W determined by physical laws. Nevertheless, we can word the problem as "this is how it ought to be, given certain ethical values", which will form W. Then we want the algorithm to follow this worldview W. Whether this worldview or the ethical values are prescribed by legal rules or could be seen as a constructed reality does not matter for the algorithm. Just as in physics, where we can say "energy is always conserved"—and if it is not there is a problem in the data and/or the ML model—, in a socio-technical setting we can say "there should not be a difference between different genders with respect to X" – and if the ML model does not adhere to this, this means there is a problem in the data and/or the ML model. Paralleling the discussion about PIML, also in this setting the data D may not be aligned with W. This may be due to selection bias (e.g., some demographic groups are overrepresented in D), but may also be caused by relationships in data that are created in conflict with W (for example recruitment datasets that were created by human recruiters, who have biases themselves). While there are clearly philosophical and practical differences between these issues, and the issues that, e.g., PIML tries to solve, it is still possible to equally describe both problems in the same language with our high-level conceptual model.

There is no single definition of fairness in general, and, therefore, also no general definition of a fair ML algorithm (see, e.g., [1] for different approaches to fairness metrics). The notion of fairness depends on context (including cultural context) and the application at hand. However, an important aspect is that different fairness constraints usually contradict each other [1]. There are also more complex approach to fairness, for example the question of long-term fairness in changing systems [5, 15, 17, 21].[9] gives an overview of work on ethical computing, specifically for software that incorporates data-driven approaches. [6] presents a framework for incorporating ethical principles in technology-enhanced learning software. IEEE is currently working on establishing a standard for ethically aligned design for autonomous and intelligent systems [4] The choice of the fairness definition in our setting is part of our worldview. The exact (quantitative) definition cannot be independent of the used ML model, and is part of translating W to C.
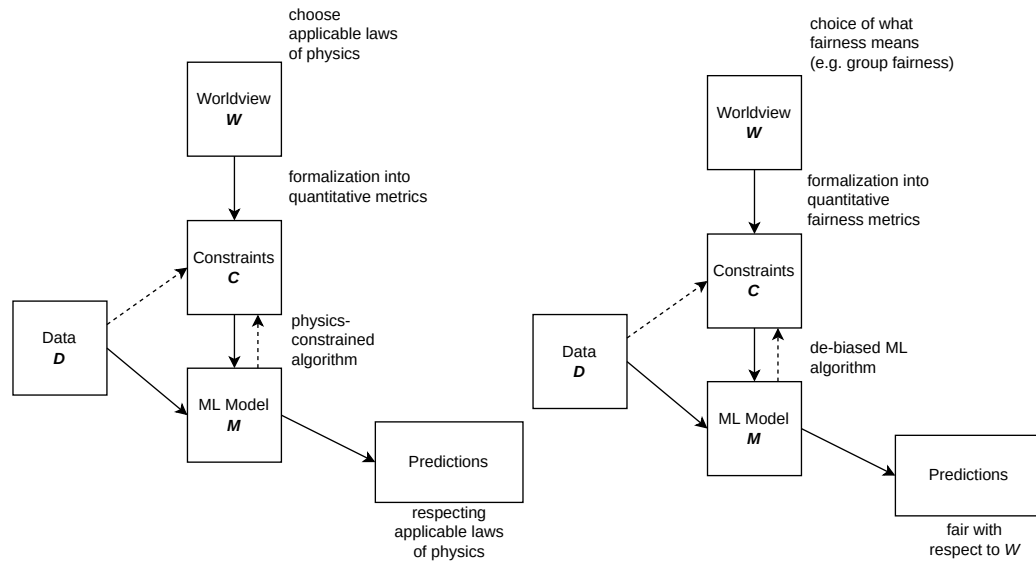
Fig. 2. Conceptual model applied to physics informed ML (right) and fair ML (left).

Many different approaches have been proposed for incorporating different fairness constraints into ML models [2, 25]. This is often called reducing bias, or "debiasing". Debiasing algorithms either work by modifying the training data, modifying the training algorithm, post-processing the outputs of the ML model, or are a combination of those approaches. In our high-level language, all these methods can be incorporated into M, while the chosen fairness definition forms W. The debiasing algorithms then use this definition of fairness as constraint. Some algorithms can work with hard constraints (ensuring that the target fairness is reached), others work with soft constraints (e.g., integrating it in the loss function, in which the training algorithms balances it with other targets such as accuracy).

In contrast to PIML, different fairness definitions usually contradict each other, and we need to *choose* which approach we take. For example, in a setting where fewer individuals from one gender than from another apply to a certain job, it is impossible to meet both the criterion that the same number of individuals of both genders are selected (which would be a valid fairness criterion), and at the same time meet the constraint that the acceptance rates must be the same for both genders (which would also be a valid fairness criterion).

Another important difference between the setting of PIML and the setting of fairness is that the "real world" can be biased, and our beliefs can thus be contradictory to the real world. Take the following example: One might want to prescribe the moral constraint that gender must not play any role in a recruitment process. If a recruitment process is trained on historical data, and recruiters in the past were biased, then the system is trained on data that stems from a process that does not fulfill the fairness constraint that we want to impose. This is not the case for physical constraints. Here of course there can also be errors in the data etc., but fundamentally the underlying system adheres to the constraints. The conceptual model that we present in this paper is still able to describe the two concepts in a single framework. As the choice of fairness approach is part of W, we can use the same conceptual model for fairness as well as for PIML. An illustration of incorporating ethical fairness requirements in our conceptual model is shown in the right panel of fig. 2.

3.3 Other domains

We have exemplarily shown how our conceptual model relates to the fields of fair ML and physics informed ML. However, the problem that models trained on D are not satisfactory for an application, even when the algorithm is appropriate for D, does occur also in many other settings. One such area is domain generalization, which describes how well an algorithm that was trained on data from one domain can make predictions on another domain [3, 8, 18]. One could argue that this is also a case where without special measures the predictions of the raw ML model are satisfactory on the training domain, but unsatisfactory on the domain it is used in the end. Depending on the exact context, this might or might not be the case: In cases where some data from the target domain is available at training time, the problem of domain generalization would be absorbed into D in our high level view, and it would be down to finding a suitable raw model M. However, in cases in which we do not have data from the target domain, but some other information that can - at least in principle - be translated into a constraint, applying our conceptual model makes sense, and might help in adapting solutions from other domains (e.g. fairness) for the specific domain generalization problem one is faced with.

Another related area of research is psychology-informed recommender systems [14]. This type of recommender system is based on the idea that human decision support should also follow guidelines of human intelligence rather than solely follow guidelines of artificial intelligence (e.g., [13, 22]). One example is the integration of a human memory model (representing W) into recommendation algorithms (representing M) to accurately model interest shifts based on the power law of time-dependent decay of item exposure in human memory (representing C) [12].

4 CONCLUSION AND OUTLOOK

Data-driven models have enabled breakthroughs and new applications in a wide range of domains. It has, however, also become apparent that in many settings, the models still make predictions that are, in one way or the other, flawed. We referred to such models and their predictions as "unsatisfactory". To improve these unsatisfactory models, many methods have been developed that help to include into the models' knowledge and constraints not present in the training data. In this paper, we have presented a conceptual model that offers a high-level view on how to incorporate additional requirements into data-driven models. We have exemplified the conceptual model on two topics, namely Fair Machine Learning (Fair ML) and Physics-Informed Machine Learning (PIML). When looked at in detail, the two areas both use different methods to tackle different problems. However, with our concept, the approaches can be described in a common language. We hope that this will help in allowing the transfer of ideas between different domains - such as between fair ML and physics informed ML. We are convinced that approaching the different domains and areas on such a high-level view will allow identifying the concepts that can be transferred and will help in solving at least some of the open challenges that stem from relying on data-centric approaches.

REFERENCES

[1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and machine learning: Limitations and opportunities*. fairmlbook.org.

[2] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John T. Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. 2018. AI fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *CoRR* abs/1810.01943 (2018). http://arxiv.org/abs/1810.01943 arXiv: 1810.01943 tex.bibsource: dblp computer science bibliography, https://dblp.org tex.biburl: https://dblp.org/rec/journals/corr/abs-1810-01943.bib tex.timestamp: Tue, 30 Oct 2018 10:49:09 +0100.

[3] Gilles Blanchard, Gyemin Lee, and Clayton Scott. 2011. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems* 24 (2011).

[4] Raja Chatila and John C. Havens. 2019. The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. In *Robotics and Well-Being*, Maria Isabel Aldinhas Ferreira, João Silva Sequeira, Gurvinder Singh Virk, Mohammad Osman Tokhi, and Endre E. Kadar (Eds.). Vol. 95. Springer International Publishing, Cham, 11–16. https://doi.org/10.1007/978-3-030-12524-0_2 Series Title: Intelligent Systems, Control and Automation: Science and Engineering.

[5] Alexander D'Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D. Sculley, and Yoni Halpern. 2020. Fairness Is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 525–534. https://doi.org/10.1145/3351095.3372878

[6] Sebastian Dennerlein, Christof Wolf-Brenner, Robert Gutounig, Stefan Schweiger, and Viktoria Pammer-Schindler. 2020. Guiding Socio-Technical Reflection of Ethical Principles in TEL Software Development: The SREP Framework. In *Addressing Global Challenges and Quality Education (Lecture Notes in Computer Science)*, Carlos Alario-Hoyos, María Jesús Rodríguez-Triana, Maren Scheffel, Inmaculada Arnedillo-Sánchez, and Sebastian Maximilian Dennerlein (Eds.). Springer International Publishing, Cham, 386–391. https://doi.org/10.1007/978-3-030-57717-9_32

[7] N. Benjamin Erichson, Michael Muehlebach, and Michael W. Mahoney. 2019. Physics-informed Autoencoders for Lyapunov-stable Fluid Flow Prediction. https://doi.org/10.48550/arXiv.1905.10866 arXiv:1905.10866 [physics].

[8] Ishaan Gulrajani and David Lopez-Paz. 2020. In Search of Lost Domain Generalization. *CoRR* abs/2007.01434 (2020). arXiv:2007.01434 https://arxiv.org/abs/2007.01434

[9] Brittany Johnson and Justin Smith. 2021. Towards Ethical Data-Driven Software: Filling the Gaps in Ethics Research & Practice. In *2021 IEEE/ACM 2nd International Workshop on Ethics in Software Engineering Research and Practice (SEthics)*. 18–25. https://doi.org/10.1109/SEthics52569.2021.00011

[10] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. 2021. Physics-informed machine learning. *Nature Reviews Physics* 3, 6 (June 2021), 422–440. https://doi.org/10.1038/s42254-021-00314-5 Number: 6 Publisher: Nature Publishing Group.

[11] Anuj Karpatne, William Watkins, Jordan Read, and Vipin Kumar. 2018. Physics-guided Neural Networks (PGNN): An Application in Lake Temperature Modeling. arXiv:1710.11431 [cs.LG].

[12] Dominik Kowald, Subhash Chandra Pujari, and Elisabeth Lex. 2017. Temporal Effects on Hashtag Reuse in Twitter: A Cognitive-Inspired Hashtag Recommendation Approach. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*.

[13] Emanuel Lacic, Dominik Kowald, Paul Seitlinger, Christoph Trattner, and Denis Parra. 2014. Recommending Items in Social Tagging Systems Using Tag and Time Information. In *1st International Workshop on Social Personalisation co-located with Hypertext'14 (SP@HT '14)*.

[14] Elisabeth Lex, Dominik Kowald, Paul Seitlinger, Thi Ngoc Trang Tran, Alexander Felfernig, Markus Schedl, et al. 2021. Psychology-informed recommender systems. *Foundations and Trends® in Information Retrieval* 15, 2 (2021), 134–242.

[15] Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2018. Delayed impact of fair machine learning. In *International Conference on Machine Learning*. PMLR, 3150–3158.

[16] Lu Lu, Raphael Pestourie, Wenjie Yao, Zhicheng Wang, Francesc Verdugo, and Steven G Johnson. 2021. Physics-informed neural networks with hard constraints for inverse design. *SIAM Journal on Scientific Computing* 43, 6 (2021), B1105–B1132.

[17] Hussein Mouzannar, Mesrob I. Ohannessian, and Nathan Srebro. 2019. From Fair Decision Making To Social Equality. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) *(FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 359–368. https://doi.org/10.1145/3287560.3287599

[18] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. 2013. Domain Generalization via Invariant Feature Representation. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28* (Atlanta, GA, USA) *(ICML'13)*. JMLR.org, I–10–I–18.

[19] M. Raissi, P. Perdikaris, and G. E. Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378 (Feb. 2019), 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

[20] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. 2020. Learning to Simulate Complex Physics with Graph Networks. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 8459–8468. https://proceedings.mlr.press/v119/sanchez-gonzalez20a.html

[21] Sebastian Scher, Simone Kopeinik, Andreas Trügler, and Dominik Kowald. 2023. Modelling the long-term fairness dynamics of data-driven targeted help on job seekers. *Scientific Reports* 13, 1 (Jan. 2023), 1727. https://doi.org/10.1038/s41598-023-28874-9

[22] Paul Seitlinger, Dominik Kowald, Simone Kopeinik, Ilire Hasani-Mavriqi, Elisabeth Lex, and Tobias Ley. 2015. Attention please! a hybrid resource recommender mimicking attention-interpretation dynamics. In *Proceedings of the 24th International Conference on World Wide Web*. 339–345.

[23] Justin Sirignano and Konstantinos Spiliopoulos. 2018. DGM: A deep learning algorithm for solving partial differential equations. *J. Comput. Phys.* 375 (Dec. 2018), 1339–1364. https://doi.org/10.1016/j.jcp.2018.08.029

[24] Philip S. Thomas, Bruno Castro da Silva, Andrew G. Barto, Stephen Giguere, Yuriy Brun, and Emma Brunskill. 2019. Preventing undesirable behavior of intelligent machines. *Science* 366, 6468 (Nov. 2019), 999–1004. https://doi.org/10.1126/science.aag3311 Publisher: American Association for the

Advancement of Science.

[25] Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *2018 ieee/acm international workshop on software fairness (fairware)*. 1–7. tex.organization: IEEE.