The Missing Interdiscipline: Reasons and Ways to Study the History of Machine Learning

Michał Grzejdziak-Zdziarski*
Marco Loog†

Abstract

We investigate and reflect upon the common grounds for writing the history of machine learning. It may seem natural for machine learning practitioners—i.e., all people involved in working with machine learning: researchers, engineers, executives, ethicists, etc.—to think of history as a mere repository of curiosities. However, through a comprehensive study of a large number of literary sources, we argue that our discipline can benefit from it in practical ways. In particular, we discuss three areas of opportunity for machine learning practitioners provided by the study of history. Concrete examples are provided in support of our views. Subsequently, we observe, however, that the state of the art of the history of machine learning does not realize its full potential to have a positive impact on the discipline. We attribute this to three major challenges that impede historical studies of machine learning. Following a description of these challenges, we present several postulates to facilitate history research. In conclusion, we call on the machine learning community to take up the quest for its history and make space for dedicated research into the history of machine learning following historiographic standards.

*mgrzejdziak@gmail.com, Nomagic, Warszawa, Poland.
†marco.loog@ru.nl, Radboud University, Nijmegen, The Netherlands

1 Introduction: setting the stage

The growth of the machine learning community and its societal impact over the last decade is rapid and evident. It manifests in many ways. For instance, the number of submissions and accepted papers in machine learning journals and conferences grew more than four-fold between 2011 and 2021 (Pedregosa et al., 2022; Audibert et al., 2022), the search interest as reported by Google Trends increased more than ten times over the past 15 years (Trends, 2023), and the machine learning market size was estimated to have risen from $1.58 billion in 2017 (Columbus, 2017) to $15.44 billion in 2021 (Insights, 2023). Moreover, machine learning applications break records in the speed of adoption, with ChatGPT reaching 100 million monthly users in two months (Hu, 2023), while low-cost alternatives like DeepSeek can catalyze unprecedented losses in market value (Carew et al., 2025). Though not certain, it is easy to imagine that these trends will continue.

The machine learning community's reactions to this growth pace are ambivalent. There are many conflicting visions of what values we should share. The determination of some to take research in particular directions is met with confusion or opposition by those who do not understand or do not agree with these tendencies. This situation leads to tensions, which can turn into open conflicts. For instance, in 2018, following the social media campaign #ProtestNIPS led by leading researchers, the major machine learning conference Neural Information Processing Systems changed its original offensive acronym to NeurIPS, despite previous social media poll results in favor of keeping the old one (Brown, 2018b,a; Else, 2018). It triggered a heated Twitter discussion between two prominent scientists Pedro Domingos and Anima Anandkumar. In 2020, it amplified and led to hundreds of reactions amongst community members, when NeurIPS introduced its ethical guidelines (Soper, 2020). For Domingos, it was "alarming that NeurIPS papers are being rejected based on 'ethics reviews'" while his opponents in the discussion were supporting the changes. Among other outcomes, the heated conflict resulted in Anandkumar removing her Twitter account (Anandkumar, 2020) and Domingos, along with many signatories, sending an open letter to the Communications of the ACM to protest against "repressive actions aimed at limiting the free and unfettered conduct of scientific research and debate" (Signatories, 2020). There were many more such affairs over the past few years with community members taking sides and posing questions related to community values and requirements for machine learning researchers. Some examples from recent years include the exchange between Yannic Kilcher, Timnit Gebru, and Yann LeCun over biases in machine learning systems (Kilcher, 2020), Jürgen Schmidhuber's critique concerning the Honda Prize for Geoffrey Hinton (Schmidhuber, 2020), and the latest controversies around Schmidhuber's work concerning history (Schmidhuber, 2022, 2023; Gros, 2023; Loosemore, 2023). As a community, however, we should not be worried that conflicts appear, but rather that we cannot find a common ground to resolve and learn from them.

Also, there is a pressure "to do something." The authors of the position paper *It's time to do something...* (Hecht et al., 2018) observe "a massive gap between the real-world impacts of computing research and the positivity with which we in the computing community tend to view our work (...) it is analogous to the medical community only writing about the benefits of a given treatment and completely ignoring the side effects, no matter how serious they are." Some of us support the efforts to close the gap. Some of us consider this gap an overstatement, like the signatories of Domingos' open letter to the ACM (Signatories, 2020). Some of us have mixed opinions (croissants, 2020): "I agree with the general conclusion that machine learning (just like any way of making decisions) involves ethics, but I disagree with the conclusion that every NeurIPS submission should have a concluding 1-2 paragraphs about the ethics of the submission." Besides, for some of us, it is still not clear why all of our work now needs to be reviewed in terms of ethics as well. Should we not only be machine learning practitioners but ethicists in addition? Similarly, if we are required to assess negative societal impact, how should we prepare ourselves to do it in the right way? Should we now be well-informed in sociology, politics, and other fields? How do we come to such interdisciplinary knowledge?

1.1 Aim and outline

The above sketches two important aspects of the current state in which the machine-learning community finds itself. On the one hand, we have no common ground to talk through and resolve conflicts within the community. On the other hand, it is not clear how to interdisciplinarily train machine learning practitioners to comply with the growing pressure to step out of their purely technical expertise. Our proposal in this essay for both these problems is to turn to the history of machine learning. If studied systematically and following the appropriate methodological rules, it can provide us with reliable information about our discipline with which we can resolve community conflicts. Such a historiography (i.e., a study of history) of machine learning can help us understand the perspectives of others, guide us in (at times, heated) debates, and contextualize the societal impact of our work. The knowledge and awareness of history can be the common ground for building the machine learning community. Rather than taking sides in fierce debates, our contribution should facilitate this process for the good of all.

Why does the history of machine learning not serve the mentioned purposes yet? As we will establish, it is no longer true that the historiography of machine learning "is virtually nonexistent" as Aaron Plasek (2016) puts it, but it is still too often neglected. In this essay, through the analysis of historiographical research and prior work in abutting fields, we demonstrate the benefits of history and the challenges that block the research in this area. Ultimately, we conclude with a number of postulates to the machine learning community. More specifically, our contribution is structured as follows. In Section 2, we first discuss the opportunities the history of machine learning can provide to machine learning practitioners. In this context, in Section 3, we review and analyze the current state of the art. We identify major challenges that are obstacles to high-quality research in the history of machine learning in Section 4. In Section 5, we present our postulates and conclude the paper.

2 Opportunities

The history of machine learning can be thought of as a disciplinary history within the history of science or more broadly the history of knowledge (Daston, 2017). In our discussion of the benefits of studying the history of machine learning, we draw inspiration from multiple sources concerning the benefits of studying the history of science or specific disciplinary histories. Among others, we consider the 2014 Kailath lecture *Let's not dumb down the history of computer science* given by Donald Knuth (2021), the earlier cited short 2016 essay *On the Cruelty of Really Writing a History of Machine Learning* by Plasek (2016), and a number of essays and interviews related to the benefits of the history of science for scientists (Maienschein, 2000; Maienschein et al., 2008; Creath, 2010; Chang, 2016; Passos Videira and Queijo Olano, 2020; Loncar, 2022).We primarily point out views of interest and reflect on their relevance to the history of machine learning. We additionally combine this with our personal experience as machine learning practitioners in academia and industry. Generally speaking, we are convinced that the study of history can be beneficial in at least the following three areas, which we refer to as community building, practical training, and synthetic thinking. We elaborate on these three areas in the remainder of this section.

2.1 Community building

By community building we mean all activities and processes that facilitate collaboration between machine learning practitioners: 1) sharing knowledge, 2) promoting good practices, appreciating the work of others, identifying and fighting long-lasting problems, and, last but not least, 3) defining the scope of the field, its ambitions and challenges.

Here the history of machine learning has many uses. The first point above, knowledge sharing, concerns such important areas as education, exchange of ideas, or research communication. Education is one of the five main uses of history for Maienschein (2000); Maienschein et al. (2008). For Knuth (2021), history storytelling is "the best way of exposition." As he puts it: "I can be a much better teacher and writer if I can understand why it is that Leibniz didn't get it." In the context of machine learning, a notable example is the technique of back-propagation for training neural networks, which nowadays we look upon as a simple application of the ancient chain rule. Neither Rosenblatt (1958) nor Ivakhnenko (1971) may have needed it and it seems to have taken decades before the relevance of back-propagation for neural networks had been fully realized by the community.

Secondly, by tracking the origins of ideas and their development over time, history can be used to promote good practices like proper credit assignment. This point, along with the one about the role of history in "celebrating the contributions of many cultures", is made by Knuth (2021). By combining all sources of knowledge, history is a way to embrace and support diversity in the community and can help us realize and fight its long-lasting problems. A concrete example is gender bias; the fact that women are underrepresented in the machine learning community. It is estimated that only 12% of the leading machine learning researchers are female

(Fatemi, 2020). The study of history can help us understand its roots and tackle such problems.

Finally, through history, we shape the boundaries of our discipline and define our common ambitions, challenges, and methods to address them. We do it by specifying the ideas and people from the past related to machine learning, this way defining our scope of interest today. For example, in this manner, we sustain the strong connection of machine learning with the field and challenges of artificial intelligence, while we keep it looser with the fields of pattern recognition or operations research.

2.2 Practical training

By practical training, we mean that history can help machine learning practitioners perform their jobs better. We illustrate this point from the two complementary perspectives of a scientist and a historian of science, respectively. The first is (Knuth, 2021) again and the second by Hasok Chang (2016), who is a renowned historian and philosopher of science.

For Knuth (2021), history helps researchers in "understanding the process of discovery," "understanding the process of failure," and "learning how to cope with life." We can give a concrete example for the first point from the personal accounts of Judea Pearl. In a short retrospective essay, Pearl (2006) gives a step-by-step overview of the process that led him to the invention of Bayesian networks, starting from his "naive" speculations about the human mind, through the realization that conditional independence is "the most crucial factor facilitating distributed computations," and ending with connecting graphs with probability theory to yield what he called Bayesian networks. This example can guide other researchers in what questions they might ask and what explorations to perform in various stages of their research. History is a way of sharing experiences of people very distant in time and space, regardless of their professions. By studying the history of the people or institutions in the past, one can learn from their experiences about how to do their job better.

Chang (2016) focuses on what he calls "internal" functions of the history of science. He differentiates two kinds of such functions: "orthodox" and "complementary." One of the orthodox functions of the history of science is that "historical knowledge can help us understand better the scientific knowledge that we accept at present." He argues that history helps scientists to be more critical and understand scientific concepts with more depth. On the complementary side, he observes that "generally speaking, learning history is a wonderful way of opening one's mind to new possibilities." The new possibilities are understood in three senses: 1) that science could have developed in different directions, 2) that the ways of thinking in the past were radically different than the ones we experience in the present and provide different perspectives, 3) that past systems of science were different. Here, a concrete example from the history of machine learning is the fact that Rosenblatt realized his perceptron in hardware rather than software (Sejnowski, 2020); a fact that many researchers may not realize anymore. Nowadays, neural network computations are performed on general-purpose computers and it is highly non-obvious that it was actually more efficient in Rosenblatt's times to have specialized hardware. Had such a situation persisted, we could have faced different challenges and used different paradigms in the field of deep learning. Lorraine Daston, a prominent historian of science, makes a similar point about such alternative possibilities (Passos Videira and Queijo Olano, 2020): "The history of science can also provide a sense of intellectual possibility. It's very important especially for younger scientists to realize that the current reigning orthodoxy in science is not without alternatives. We did not always think the way we do now and in all probability will think differently in the future.".

2.3 Synthetic thinking

Synthetic thinking, positioning one's work in context, is indispensable in the daily work of all machine learning practitioners. Having an overview of what is going on in the field is equally important for researchers, engineers, and others. History, as a discipline combining multiple sources to construct consistent narratives, can help individuals get a wider perspective on the field.

In the most general case, history is about the "integration of many different strands of evidence" as Daston puts it in (Loncar, 2022). The most common examples in the machine learning community are "related work" sections, part of virtually all peer-reviewed papers published in journals and conferences, or literature surveys (Creath, 2010). Good literature reviews help researchers position their work in the context of the work of other researchers who addressed similar challenges or applied similar methods. Beyond literature surveys, history is a way to have a bird's eye view of the field, which facilitates specifying one's goals and leads to novel research directions or application areas. This kind of orientation that scientists may seek in the history of science is also mentioned by Daston (Passos Videira and Queijo Olano, 2020).

In addition, history can help machine learning practitioners reason about the societal impact of their work and, in Daston's words (Passos Videira and Queijo Olano, 2020), it helps "reflect on their social and political responsibility by providing case studies of how even the best of intentions have sometimes resulted in terrible human tragedies." In the context of machine learning, this lesson from history is as important as in any other discipline. In (Passos Videira and Queijo Olano, 2020), Daston further observed that "scientists often take their own categories from the colloquial categories of the society in which they were raised, – how could it be otherwise? – without further critical reflection." For example, this problem concerns the collection of large datasets that are core to machine learning progress. In the famous ImageNet computer vision dataset there were images with such categories as "call girl", "bad person", or "hypocrite" (Crawford and Paglen, 2021). Though these categories were not used for training current state-of-the-art computer vision models, the images of the actual people labeled with these categories were available publicly on the Internet for a long time. It is obvious now that our datasets are biased and we have to be aware of these biases and how to control them. As Plasek (2016) suggests: "to begin to see how biases are propagated and reinforced via machine learning system training data" we need history or, specifically, "histories of datasets".

3 State of the art

In the previous section, we explained that the study of history can be beneficial for machine learning practitioners in many ways. In this section, we review the state of the art and analyze whether it provides the mentioned opportunities. We performed our literature survey by querying internet search engines, consulting scientific publication databases, and following references in already surveyed papers. We took into account works that are directly related to machine learning and, as their main topic, are concerned with something located in the past, with the development of something over time, or focus on history in itself. The result of our survey is more than 40 works, which we grouped into six categories concerning their methodology and genre to assess their relevance in light of the benefits listed in the previous section. The categories we defined, and subsequently discuss, are: 1) popular histories, 2) reflections, 3) literature surveys, 4) analytical historiography, 5) philosophical historiography, 6) socio-political historiography.[1]

3.1 Popular histories

Popular histories are personal accounts about certain historical topics given usually by machine learning practitioners themselves based on their own experiences and observations. They are usually written to be good histories to read without a strict methodology. Characteristic features of popular histories are overexposure of the past directly related to the author, the use of personal anecdotes and correspondence, focus on technical and financial aspects, explicit contribution to a specific worldview, and subjective interpretations. The classical examples of popular histories are two extensive books about the history of artificial intelligence in which machine learning forms a large part by Crevier (1993) and Nilsson (2009). There are also a book and a paper about the history of deep learning by Sejnowski (2018, 2020). Other related examples are Cordeschi (2007); Tappert (2019); Fradkov (2020); Stanko (2020); Adami (2021); Foote (2022). To this category, we count also memoirs and interviews, in particular the two series of interviews by Anderson and Rosenfeld (1998) and Ford (2018). Finally, somewhat related to this group are the historical essays in a book by Lungarella et al. (2007), but they are concerned mostly with artificial intelligence rather than machine learning.

Popular histories contribute to all three areas of community building, practical training, and synthetic thinking. By giving an overall perspective on the field and its ambitions, they help shape the boundaries of what in the past was thought of as machine learning. Their storytelling assets, the use of anecdotes, and subjective

[1] Roughly, we came to the proposed categorization in the following way. If the work's authors and/or venue are from humanities or philosophy, it approximately identifies with the philosophical historiography. If coming from the social sciences the work typically identifies with the socio-political historiography. When from mathematics or the natural sciences a further distinction in four categories applies. If the work is peer-reviewed and has an underlying research question it typically categorizes as analytic historiography and as literature research if it doesn't pursue a particular research question. On the other hand, non-peer-reviewed work roughly identifies with reflections if they try to argue a particular case and with popular histories if they don't.

but strong opinions can provide not only educational help in engaging students, but also inspiration for one's work. Many popular histories cover periods of decades and a wide range of areas, giving this way a good overview of the methods of machine learning developed over time. On the other hand, the danger of popular histories is their bias, which, unintentionally, can lead to a widespread of false or even harmful claims. This can, in turn, perpetuate rather than fight long-lasting problems in the community. Moreover, their focus on technical aspects goes hand in hand with a poor treatment of societal issues.

3.2 Reflections

Reflections are in-depth essays of machine learning practitioners about the field and its history. Their main goal is on a meta-level: rather than presenting and analyzing specific facts, they focus on naming topics for further analysis. An example is a paper by Newell (1983) which identifies the intellectual issues in the history of artificial intelligence. A paper by Breiman (2001) names two paradigms of approaching statistical data modeling which gave way to two separate research communities. A retrospective by Pearl (2006) is a reflection on the research process which led towards a major contribution. Marcus (2018) evaluates the recent history of deep learning to reflect on ten concerns of this subdiscipline of machine learning. There is also Chauvet (2018) who suggests that every 30 years the debate in AI repeats some patterns. To this category, we can also count the summary of the NeurIPS 2019 Retrospectives Workshop by Sodhani et al. (2020). Finally, a paper by Hooker (2020) is a reflection on the influence of the availability of specific hardware at a specific time on research progress.

Reflections, due to their unrestricted style and wide range of tackled topics, can contribute to any community building, practical training, and synthetic thinking. They can trigger discussions whose results can be beneficial for the community in the long perspective. They can point to specific processes that impact the work of individuals, giving them tools to situate their work in context. It must be noted that often they are not peer-reviewed and are published as technical reports or preprints. The consequence of this is that the author's biases might not be identified and addressed properly in the review process. Also, for the same reason the impact of reflections is limited.

3.3 Literature surveys

Literature surveys are works that briefly summarize developments in machine learning in specific topics or in general, using extensive numbers of technical sources. Short and focused literature surveys are parts of most machine learning papers in the form of "related work" sections, but there are also works that are dedicated to a profound review of some topics. Probably the most prominent examples of such standalone literature surveys are two papers by Schmidhuber (2015, 2022). Schmidhuber (2015) concisely reviews papers that contributed to fundamental ideas in deep learning in neural networks. As the author admits, the paper "mostly consists of references". A follow-up paper by Schmidhuber (2022) contains more discussions and a wider thematic scope, but it still concerns more than five hundred references on less than thirty pages. Another example of a literature survey is a review by Carbonell et al. (1983) that briefly discusses methodological approaches to machine learning, its objectives, and historical development. The authors propose different taxonomies of learning systems but do not treat them in detail. Other examples of literature surveys are Rosenfeld and Wechsler (2000); Misra and Saha (2010); Bengio et al. (2012); LeCun (2019); Loog et al. (2020); Michelucci (2024).

Due to their focus on tracking the specific topic over time, extensive literature surveys are usually the most reliable way for proper credit assignments. Due to their brevity and limited depth of treatment, they can provide only high-level overviews and starting points for researchers interested in specific topics. The characteristic of literature surveys is that they are rigorous, but their ambitions are limited. As a result, except for some specific aspects, they contribute relatively little to the aspects of community building, practical training, and synthetic thinking.

3.4 Analytical historiography

Analytical historiography is the study of the past in terms of time-series analysis. The main method is to reduce certain historical events or processes to numbers and analyze the evolution of these numbers over time. This genre, like the three above, is again written by machine learning practitioners who apply the methods of machine learning to analyze history. There are extensive studies very much based on numeric data and measurements like Cardon et al. (2018); Gurcan and Sevik (2019); Ahmed and Wahed (2020); Audibert et al. (2022). Another example is the recent short JMLR retrospective blog post by Pedregosa et al. (2022) which showcases plots of number-reduced characteristics of the JMLR journal over 20 years of its existence with little interpretation.

Analytical historiography can be valuable for machine learning practitioners in the sense that it can help identify trends in the data that can highlight certain problems or challenges in the field, but it usually does not provide tools to understand the reasons for such trends. For this, one needs to go beyond quantitative analysis.

3.5 Philosophical historiography

Philosophical historiography is the kind of writing history that is performed mainly by philosophers or machine learning practitioners with a philosophical background. It concerns relations between philosophy and machine learning, in particular, how these two influence each other. A book by Mackenzie (2017) is a philosophical attempt to express machine learning within the framework of the archaeology of knowledge developed by Foucault (1972). Berkeley (2019) is a review of philosophical issues which were inspired by three waves of connectionism. How the emergence of specific machine learning models is related to value shifts is discussed by Dotan and Milli (2019). Another notable work is the PhD thesis of Grimsley (2022).

Philosophical historiography, with its intellectual depth, can be an inspiring tool for machine learning practitioners to understand and identify the sources and implications of their work in a contemporary intellectual context. This way it can contribute to building depth of machine learning arguments, in particular in their relation to the question of what intelligence is. However, it might be quite inaccessible due to its grounding in philosophical traditions unknown to most machine learning practitioners. Therefore, its contributions might be limited.

3.6 Socio-political historiography

Socio-political historiography is authored mostly by historians, sociologists, and ethicists and concerns a socio-political context of machine learning practice. The works in this category started appearing in particular in the last decade following the increase in the societal impact of machine learning after deep learning breakthroughs in computer vision and natural language processing. The only work that we found in this category from before 2012 is a paper by Olazaran (1996) which analyzes the seminal book by Minsky and Papert (1969) within the framework of the sociology of scientific knowledge. After 2012 many papers were devoted to machine learning datasets, and how they are used and constructed. These papers concern mostly very recent history so many of them can be treated as both historical and sociological. Crawford and Paglen (2021) adopt Foucault's *archaeology* (Foucault, 1972) to study popular computer vision datasets. Denton et al. (2020) and Denton et al. (2021) adopt another framework by Foucault, so-called *genealogy* (Foucault, 1979). Other papers concerning datasets and their use are Jo and Gebru (2020); Paullada et al. (2021); Koch et al. (2021). Another kind of work in this category concerns politics or economy from a wider perspective. Examples are Mohamed et al. (2020); Crawford (2021); Penn (2021); Steinhoff (2021)

The studies of socio-political historiography can provide machine learning practitioners with the knowledge to reason about the societal impact of their work. As the authors of the works in this category are trained specialists in social sciences and historiography, they are usually more critical and less biased than technical authors. If applied to the history of the machine learning community, socio-political historiography can help identify and address its long-lasting problems. On the other hand, due to its non-technical focus, it often fails to attract machine learning practitioners.[2] In part, the feeling is that it does not offer insight that can be directly applied to their daily technical work.

4 Challenges

The previous section substantiates that the state of the art in the historiography of machine learning is largely biased (popular histories and reflections), limited (literature surveys and analytical historiography), or unattractive for machine learning practitioners (philosophical historiography and socio-political historiography). Can we have a historiography without all these issues? In this section, we identify and

[2] We hope, of course, that the current work is spared such fate.

describe three challenges that hamper establishing a good, unbiased, ambitious, and attractive history of machine learning: 1) lack of communication with historians, 2) confusion about the topic, 3) contemporaneity.

4.1 Lack of communication with historians

History is the domain of historiography, an academic discipline in itself with its own community and publishing venues. In our discussion in Section 3, we observed that historians contribute to the history of machine learning mostly by practicing what we called socio-political historiography. It focuses on aspects of what we identified in Section 2 as community building and synthetic thinking, but does not necessarily offer much in the area of practical training. Probably for this reason, its reception in the machine learning community is limited.

The problem of lack of communication between historians and practitioners concerns many contemporary disciplines, as has been frequently mentioned by historians of science. Daston (2009) observes that "scientists stopped reading the history of science" because "it had succeeded all too well in making past science wholly unfamiliar". Maienschein et al. (2008) note that "scientists usually do not assume that the history and philosophy of science could be of value for their actual research". This problem manifested itself in two recent discussions between scientists and historians of science in mathematics and computer science, respectively.

In 2014, Viktor Blåsjö (2014) and Michael Fried (2014), both established historians of mathematics, disputed the "modern consensus in the historiography of mathematics" in the *Journal of Humanistic Mathematics*. Blåsjö (2014) criticized a large body of modern historiography of mathematics as practiced by professional historians. His main allegation was that it focused on keeping high historiographic standards but failed to be attractive or relevant for audiences other than historians themselves. In particular, it wouldn't appeal to mathematicians, mathematics students, and teachers. In response, Fried (2014) argued that the work of historians is to understand the past in its own right, with as little impact from the present as possible. For him, the kind of history of mathematics that Blåsjö (2014) postulated was mathematics rather than history. As such, it should rather be written by mathematicians themselves.

An analogous exchange took place on the pages of *IEEE Annals of the History of Computing* from 2007 to 2014 between Martin Campbell-Kelly (2007), Donald Knuth (2021), and Thomas Haigh (2015). In his reflective overview of the historiography of computing, Campbell-Kelly (2007) observed a switch in focus from technical matters to socio-political issues of computing and assessed this change as clearly desirable. Knuth (2021) admitted that socio-political issues were also important, but resigning from studying technical aspects at all was limiting the value of historiography for computer scientists and more technical audiences. Haigh (2015) responded to Knuth that the kind of history that he postulates can only be practiced if computer scientists get involved in writing history themselves. He suggested that for it to happen, computer scientists should build institutional context with dedicated funding, tenure positions, PhD studies, and publishing venues.

Looking at the reviewed literature, we think that the lack of communication with historians is the main reason behind the three problems we sketched at the beginning of this section. The bias of reflections and popular histories and the limited value of literature surveys and analytical historiography arise from the fact that machine learning practitioners have no training in historiographic methodology. Without communication with historians, common publishing venues, or interdisciplinary teams, they hardly have ways or incentives to receive such training. On the other hand, historians who write about the history of machine learning, do it mostly for other audiences than the machine learning community and have no incentives to make it attractive for machine learning practitioners.

4.2 Confusion about the topic

The strong relation between machine learning and artificial intelligence combined with little methodological reflection results in blurring the subject matter of the historiography of machine learning. Authors of historical works about machine learning focus either on machine learning and its methods[3] or on artificial intelligence, treating machine learning merely as one of its approaches, sometimes even omitting any explicit reference to machine learning[4]. This concerns both the authors with more technical and more sociological background.

There are two problems with this split. First, it results in additional fragmentation of the historiography of machine learning into two groups. This adds to research communication problems that are consequences of a lack of communication between scientists and historians. Second, the focus of one group on artificial intelligence makes it oblivious to important aspects specific to machine learning which limits its value. The problem was discussed by Plasek (2016). The author observes that the focus on the ill-defined concept of artificial intelligence made historians neglect the construction and maintenance of large annotated datasets, which is an inherent part of machine learning practice, but not of other approaches to artificial intelligence. As Plasek (2016) argues, to understand the societal impact of machine learning, it is necessary to study the histories of datasets, because they encode and perpetuate human biases which are later included in machine learning algorithms.

4.3 Contemporaneity

The history of machine learning is necessarily a kind of contemporary history. This poses a number of additional challenges. The work of historians of machine learning will be incomplete as they will have limited access to information that is still classified

[3]Examples are Alom et al. (2018); Anderson and Rosenfeld (1998); Carbonell et al. (1983); Denton et al. (2021, 2020); Dotan and Milli (2019); Foote (2022); Fradkov (2020); Koch et al. (2021); Loog et al. (2020); Mackenzie (2017); Misra and Saha (2010); Olazaran (1996); Paullada et al. (2021); Plasek (2016); Schmidhuber (2015); Sejnowski (2018); Sodhani et al. (2020); Stanko (2020); Tappert (2019)

[4]Some examples are the following Adami (2021); Ahmed and Wahed (2020); Audibert et al. (2022); Chauvet (2018); Cordeschi (2007); Crawford (2021); Crawford and Paglen (2021); Crevier (1993); Grimsley (2022); Lungarella et al. (2007); Mohamed et al. (2020); Nilsson (2009); Penn (2021); Schmidhuber (2022); Sejnowski (2020); Steinhoff (2021)

for security or commercial reasons. Another challenge imposed by contemporaneity is the danger of bias, as no distance in time can make it difficult to study some topics with the necessary objectivity. Tensions may arise between an author's role as a historian and as a participant in contemporary discussions. Often, the authors themselves will be actors of history and they may be biased towards their own agenda.

The problem of contemporaneity appeared recently in several discussions on the Connectionists mailing list around Jürgen Schmidhuber's and Terrence Sejnowski's works concerning the history of deep learning. Some machine learning researchers participating in discussions discouraged studying the history of machine learning because of its contemporary character. For one participant, Schimhuber was "part of history, not a historian" (Hanson, 2022). Another participant suggested that "all these threads [of Schmidhuber's work] will be sorted out by historians one hundred years from now" (Sejnowski, 2022). Such comments can be discouraging to (prospective) authors and hamper studies in the history of machine learning.

5 Postulates and Conclusion

We argued in Section 2 that the history of machine learning can be beneficial for the practitioner. In Section 3, however, we provided strong arguments that the current state of the art is generally biased, limited, or unattractive to the machine learner. The three major obstacles in the way of historical studies of machine learning not suffering from these issues we then discussed in Section 4.

In light of the preceding reflections, and despite the challenges identified, we are convinced that studying the history of machine learning can be an effective and marked way to build our community sustainably and responsibly. Here, we come to three postulates that, if realized, could lay the foundation for a historiography of machine learning that would be methodologically solid, consistent, and practical for the machine learning community. They are as follows.

1. **Make space for history in the machine learning community.** The lack of communication with historians will be a major obstacle as long as we will expect historians to write history for the machine-learning community without giving them any incentives. But, as the examples from the disciplines of computer science and mathematics demonstrate, historians have their own institutional context and they will not all magically turn to study the history of machine learning. To facilitate such research, we need to act in two directions. First, machine learning practitioners should get basic historical training and go beyond literature surveys, popular histories, and analytical historiography. Second, the machine learning community should make space for professional historians of machine learning via incentives for historical research, for example by organizing history tracks at conferences or funding historical research. A good example of such a space for historians in computing is ACM SIGGRAPH, which sustains a history community (SIGGRAPH, Accessed on 07.05.2024a) and welcomes submissions of historical works for its yearly conference about computer graphics (SIGGRAPH, Accessed on 07.05.2024b).

2. **Refer to machine learning explicitly.** Blurring the lines between machine learning and artificial intelligence results in additional obstacles in the research communication. The focus on artificial intelligence can result in the omission of some crucial aspects of machine learning. We call historians to always refer to machine learning explicitly to facilitate knowledge exchange and limit the risk of involuntary neglect of important technical details.

3. **Do not treat contemporaneity as a barrier.** It is not possible to avoid the challenges related to the contemporaneity of machine learning when writing its history. Though some argue that we should wait decades or hundreds of years for historians to treat it, it is not clear if machine learning will ever stop being contemporary.

We are aware that realizing these postulates is not straightforward. It is not a task for a single person or a small research team. It is a continuous process; at no point in time will we be able to say that the history of machine learning is complete. Therefore, contributions in all shapes and sizes ought to be welcomed. It can be the monthly reading of a history paper and discussing it with colleagues. It can be developing and teaching a course on the history of machine learning or opening up academic positions, starting with postdoc positions and PhD studies. Members of editorial boards and conference chairs can invite history contributions to their venues and those who want to delve even deeper into our machine learning history may choose to get methodological historiographic training and start writing themselves. We believe every member of the machine learning community can find a fitting role in this venture.

With the rapid pace of the field, we need to start building foundations for a professional historiography of machine learning. This will not happen by itself, and we call on the machine learning community to actively take up the quest for the history of machine learning. Big or small, it would be great to have your contribution as well.

References

Christoph Adami. A Brief History of Artificial Intelligence Research. *Artificial Life*, 27(2):131–137, 2021. doi: 10.1162/artl_a_00349. URL https://doi.org/10.1162/artl_a_00349.

Nur Ahmed and Muntasir Wahed. The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research, 2020. URL https://arxiv.org/abs/2010.15581.

Md Zahangir Alom, Tarek M. Taha, Christopher Yakopcic, Stefan Westberg, Paheding Sidike, Mst Shamima Nasrin, Brian C Van Esesn, Abdul A S. Awwal, and Vijayan K. Asari. The History Began from AlexNet: A Comprehensive Survey on Deep Learning Approaches, 2018. URL https://arxiv.org/abs/1803.01164.

Anima Anandkumar. My heartfelt apology. https://anima-ai.org/2020/12/16/my-heartfelt-apology/, 2020. Accessed on 09.03.2023.

James A Anderson and Edward Rosenfeld, editors. *Talking nets: An Oral History of Neural Networks*. Bradford Books. MIT Press, London, England, July 1998.

Rafael B. Audibert, Henrique Lemos, Pedro Avelar, Anderson R. Tavares, and Luís C. Lamb. On the Evolution of A.I. and Machine Learning: Towards Measuring and Understanding Impact, Influence, and Leadership at Premier A.I. Conferences, 2022. URL https://arxiv.org/abs/2205.13131.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation Learning: A Review and New Perspectives, 2012. URL https://arxiv.org/abs/1206.5538.

Istvan S. N. Berkeley. The Curious Case of Connectionism. *Open Philosophy*, 2(1): 190–205, August 2019. doi: 10.1515/opphil-2019-0018. URL https://doi.org/10.1515/opphil-2019-0018.

Viktor Blåsjö. A Critique of the Modern Consensus in the Historiography of Mathematics. *Journal of Humanistic Mathematics*, 4(2):[113]–123, July 2014. doi: 10.5642/jhummath.201402.12. URL https://doi.org/10.5642/jhummath.201402.12.

Leo Breiman. Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3): 199–215, 2001. ISSN 08834237. URL http://www.jstor.org/stable/2676681.

Jennings Brown. 'NIPS' AI Conference Changes Name Following Protests Over Gross Acronym. https://gizmodo.com/nips-ai-conference-changes-name-following-protests-ov-1830548185, 2018a. Accessed on 09.03.2023.

Jennings Brown. AI Community Feuds Over 'NIPS,' the Controversial Name of Its Top Conference. https://gizmodo.com/ai-community-feuds-over-nips-the-controversial-name-of-1830028688, 2018b. Accessed on 09.03.2023.

Martin Campbell-Kelly. The History of the History of Software. *IEEE Annals of the History of Computing*, 29(4):40–51, 2007. doi: 10.1109/MAHC.2007.4407444.

Jaime G. Carbonell, Ryszard S. Michalski, and Tom M. Mitchell. Machine Learning: A Historical and Methodological Analysis. *AI Magazine*, 4(3):69, Sep. 1983. doi: 10.1609/aimag.v4i3.406. URL https://ojs.aaai.org/index.php/aimagazine/article/view/406.

Dominique Cardon, Jean-Philippe Cointet, and Antoine Mazieres. Neurons spike back: The invention of inductive machines and the artificial intelligence controversy. *Reseaux*, 36:173–220, 2018.

Sinéad Carew, Amanda Cooper, and Ankur Banerjee. DeepSeek sparks AI stock selloff; Nvidia posts record market-cap loss. https://www.reuters.com/technology/chinas-deepseek-sets-off-ai-market-rout-2025-01-27/, 2025. Accessed on 03.03.2025.

Hasok Chang. Who cares about the history of science? *Notes and Records: the Royal Society Journal of the History of Science*, 71(1):91–107, October 2016. doi: 10.1098/rsnr.2016.0042. URL https://doi.org/10.1098/rsnr.2016.0042.

Jean-Marie Chauvet. The 30-Year Cycle In The AI Debate, 2018. URL https://arxiv.org/abs/1810.04053.

Louis Columbus. Roundup Of Machine Learning Forecasts And Market Estimates, 2020. https://www.forbes.com/sites/louiscolumbus/2020/01/19/roundup-of-machine-learning-forecasts-and-market-estimates-2020/, 2017. Accessed on 26.02.2023.

Roberto Cordeschi. AI turns fifty: revisiting its origins. *Applied Artificial Intelligence*, 21(4-5):259–279, April 2007. doi: 10.1080/08839510701252304. URL https://doi.org/10.1080/08839510701252304.

Kate Crawford. *The Atlas of AI*. Yale University Press, April 2021. doi: 10.2307/j.ctv1ghv45t. URL https://doi.org/10.2307/j.ctv1ghv45t.

Kate Crawford and Trevor Paglen. Excavating ai: the politics of images in machine learning training sets. *AI & SOCIETY*, June 2021. doi: 10.1007/s00146-021-01162-8. URL https://doi.org/10.1007/s00146-021-01162-8.

Richard Creath. The Role of History in Science. *Journal of the History of Biology*, 43(2):207–214, 2010. ISSN 00225010, 15730387. URL http://www.jstor.org/stable/40802740.

Daniel Crevier. *AI: The Tumultuous History of the Search for Artificial Intelligence*. Basic Books, Inc., USA, 1993. ISBN 0465029973.

croissants. An Open Letter to the Communications of the ACM. https://news.ycombinator.com/item?id=25575321, 2020. Accessed on 10.03.2023.

Lorraine Daston. Science Studies and the History of Science. *Critical Inquiry*, 35(4):798–813, January 2009. doi: 10.1086/599584. URL https://doi.org/10.1086/599584.

Lorraine Daston. The History of Science and the History of Knowledge. *KNOW: A Journal on the Formation of Knowledge*, 1(1):131–154, March 2017. doi: 10.1086/691678. URL https://doi.org/10.1086/691678.

Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, Hilary Nicole, and Morgan Klaus Scheuerman. Bringing the People Back In: Contesting Benchmark Machine Learning Datasets, 2020. URL https://arxiv.org/abs/2007.07399.

Emily Denton, Alex Hanna, Razvan Amironesei, Andrew Smart, and Hilary Nicole. On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, 8(2):205395172110359, July 2021. doi: 10.1177/20539517211035955. URL https://doi.org/10.1177/20539517211035955.

Ravit Dotan and Smitha Milli. Value-laden Disciplinary Shifts in Machine Learning, 2019. URL https://arxiv.org/abs/1912.01172.

Holly Else. AI conference widely known as 'NIPS' changes its controversial acronym. *Nature*, November 2018. doi: 10.1038/d41586-018-07476-w. URL https://doi.org/10.1038/d41586-018-07476-w.

Falon Fatemi. Bridging The Gender Gap In AI. https://www.forbes.com/sites/falonfatemi/2020/02/17/bridging-the-gender-gap-in-ai/, 2020.

Keith D. Foote. *The History of Machine Learning and Its Convergent Trajectory Towards AI*, chapter 7, pages 129–142. John Wiley & Sons, Ltd, 2022. ISBN 9781119815075. doi: https://doi.org/10.1002/9781119815075.ch9. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119815075.ch9.

Martin Ford. *Architects of Intelligence*. Packt Publishing, Birmingham, England, November 2018.

Michel Foucault. *Archaeology of knowledge*. Tavistock Publications, London, England, 1 edition, 1972.

Michel Foucault. *Discipline and Punish: The birth of the prison*. Vintage Books, 1979. ISBN 9780394727677. URL https://www.bibsonomy.org/bibtex/2d402f96433b9f3808bee0bd49464cde9/rlipp.

Alexander L. Fradkov. Early History of Machine Learning. *IFAC-PapersOnLine*, 53 (2):1385–1390, 2020. doi: 10.1016/j.ifacol.2020.12.1888. URL https://doi.org/10.1016/j.ifacol.2020.12.1888.

Michael Fried. The Discipline of History and the "Modern Consensus in the Historiography of Mathematics". *Journal of Humanistic Mathematics*, 4(2):[124]–136, July 2014. doi: 10.5642/jhummath.201402.13. URL https://doi.org/10.5642/jhummath.201402.13.

Christopher Grimsley. *Contextualizing Artificial Intelligence: The History, Values, and Epistemology of Technology in the Philosophy of Science*. PhD thesis, University of Kentucky, 2022. URL https://uknowledge.uky.edu/philosophy_etds/34/.

Claudius Gros. Connectionists: ?==?utf-8?q? Annotated History of Modern AI and Deep Learning. https://mailman.srv.cs.cmu.edu/pipermail/connectionists/2023-January/039227.html, 2023. Accessed on 09.03.2023.

Fatih Gurcan and Setenay Sevik. Mapping the Research Landscape of Deep Learning from 2001 to 2019. In *2019 1st International Informatics and Software Engineering Conference (UBMYK)*. IEEE, November 2019. doi: 10.1109/ubmyk48245.2019.8965595. URL https://doi.org/10.1109/ubmyk48245.2019.8965595.

Thomas Haigh. The tears of Donald Knuth. *Communications of the ACM*, 58(1): 40–44, December 2015. doi: 10.1145/2688497. URL https://doi.org/10.1145/2688497.

Stephen Jose Hanson. Connectionists: Scientific Integrity, the 2021 Turing Lecture, etc. https://mailman.srv.cs.cmu.edu/pipermail/connectionists/2022-June/038299.html, 2022. Accessed on 07.05.2023.

B. Hecht, L. Wilcox, J. P. Bigham, J. Schoning, E. Hoque, J. Ernst, Y. Bisk, L. De Russis, L. Yarosh, B. Anjam, D. Contractor, and C. Wu. It's Time to Do Something: Mitigating the Negative Impacts of Computing Through a Change to the Peer Review Process, 2018.

Sara Hooker. The Hardware Lottery, 2020. URL https://arxiv.org/abs/2009.06489.

Krystal Hu. ChatGPT sets record for fastest-growing user base - analyst note. https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/, 2023. Accessed on 26.11.2024.

Fortune Business Insights. Machine learning market size, share, growth & trends [2029]. https://www.fortunebusinessinsights.com/machine-learning-market-102226, 2023. Accessed on 26.02.2023.

A. G. Ivakhnenko. Polynomial theory of complex systems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-1(4):364–378, 1971. doi: 10.1109/TSMC.1971.4308320.

Eun Seo Jo and Timnit Gebru. Lessons from Archives: Strategies for Collecting Sociocultural Data in Machine Learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, page 306–316, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372829. URL https://doi.org/10.1145/3351095.3372829.

Yannic Kilcher. [Drama] Yann LeCun against Twitter on Dataset Bias. https://www.youtube.com/watch?v=n1SXlK5rhR8, 2020. Accessed on 09.03.2023.

Donald E. Knuth. Let's not dumb down the history of computer science. *Communications of the ACM*, 64(2):33–35, January 2021. doi: 10.1145/3442377. URL https://doi.org/10.1145/3442377. Edited by Len Shustek.

Bernard Koch, Emily Denton, Alex Hanna, and Jacob Gates Foster. Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. URL https://openreview.net/forum?id=zNQBIBKJRkd.

Yann LeCun. 1.1 Deep Learning Hardware: Past, Present, and Future. In *2019 IEEE International Solid- State Circuits Conference - (ISSCC)*. IEEE, February 2019.

doi: 10.1109/isscc.2019.8662396. URL https://doi.org/10.1109/isscc.2019.8662396.

Samuel Loncar. Does Science Need History? A Conversation with Lorraine Daston. *MRB Interviews*, 2022.

Marco Loog, Tom Viering, Alexander Mey, Jesse H. Krijthe, and David M. J. Tax. A brief prehistory of double descent. *Proceedings of the National Academy of Sciences*, 117(20):10625–10626, May 2020. doi: 10.1073/pnas.2001875117. URL https://doi.org/10.1073/pnas.2001875117.

Richard Loosemore. Connectionists: Annotated History of Modern AI and Deep Learning. https://mailman.srv.cs.cmu.edu/pipermail/connectionists/2023-January/039220.html, 2023. Accessed on 07.05.2023.

Max Lungarella, Fumiya Iida, Josh Bongard, and Rolf Pfeifer, editors. *50 Years of Artificial Intelligence*. Springer Berlin Heidelberg, 2007. doi: 10.1007/978-3-540-77296-5. URL https://doi.org/10.1007/978-3-540-77296-5.

Adrian Mackenzie. *Machine Learners*. The MIT Press, 2017. doi: 10.7551/mitpress/10302.001.0001. URL https://doi.org/10.7551/mitpress/10302.001.0001.

Jane Maienschein. Why Study History for Science? *Biology & Philosophy*, 15(3): 339–348, June 2000. doi: 10.1023/a:1006733114136. URL https://doi.org/10.1023/a:1006733114136.

Jane Maienschein, Manfred Laubichler, and Andrea Loettgers. How Can History of Science Matter to Scientists? *Isis*, 99(2):341–349, June 2008. doi: 10.1086/588692. URL https://doi.org/10.1086/588692.

Gary Marcus. Deep Learning: A Critical Appraisal, 2018. URL https://arxiv.org/abs/1801.00631.

Umberto Michelucci. Machine learning: History and terminology. In *Fundamental Mathematical Concepts for Machine Learning in Science*, pages 9–20. Springer, 2024.

Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, USA, 1969.

Janardan Misra and Indranil Saha. Artificial neural networks in hardware: A survey of two decades of progress. *Neurocomputing*, 74(1-3):239–255, December 2010. doi: 10.1016/j.neucom.2010.03.021. URL https://doi.org/10.1016/j.neucom.2010.03.021.

Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology*, 33(4):659–684, July 2020. doi: 10.1007/s13347-020-00405-8. URL https://doi.org/10.1007/s13347-020-00405-8.

Allen Newell. Intellectual Issues in the History of Artificial Intelligence. In *The Study of Information: Interdisciplinary Messages*, page 187–294. John Wiley & Sons, Inc., USA, 1983. ISBN 047188717X.

Nils J. Nilsson. *The Quest for Artificial Intelligence*. Cambridge University Press, October 2009. doi: 10.1017/cbo9780511819346. URL https://doi.org/10.1017/cbo9780511819346.

Mikel Olazaran. A Sociological Study of the Official History of the Perceptrons Controversy. *Social Studies of Science*, 26(3):611–659, 1996. ISSN 03063127. URL http://www.jstor.org/stable/285702.

Antonio Augusto Passos Videira and Juan Andrés Queijo Olano. The Power of History: an interview with Lorraine Daston. *Contemporánea*, 10(1):197–202, jul. 2020. URL https://ojs.fhce.edu.uy/index.php/cont/article/view/662.

Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11):100336, November 2021. doi: 10.1016/j.patter.2021.100336. URL https://doi.org/10.1016/j.patter.2021.100336.

Judea Pearl. Two journeys into human reasoning. Technical Report R-331, University of California, June 2006.

Fabian Pedregosa, Tegan Maharaj, Alp Kucukelbir, Rajarshi Das, Valentina Borghesani, Francis Bach, David Blei, and Bernhard Schölkopf. Retrospectives from 20 Years of JMLR. https://jmlr.org/news/2022/retrospectives.html, 2022. URL https://jmlr.org/news/2022/retrospectives.html.

Jonathan Penn. *Inventing Intelligence: On the History of Complex Information Processing and Artificial Intelligence in the United States in the Mid-Twentieth Century*. PhD thesis, University of Cambridge, 2021. URL https://www.repository.cam.ac.uk/handle/1810/315976.

Aaron Plasek. On the Cruelty of Really Writing a History of Machine Learning. *IEEE Annals of the History of Computing*, 38(4):6–8, October 2016. doi: 10.1109/mahc.2016.43. URL https://doi.org/10.1109/mahc.2016.43.

F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958. doi: 10.1037/h0042519. URL https://doi.org/10.1037/h0042519.

Azriel Rosenfeld and Harry Wechsler. Pattern recognition: Historical perspective and future directions. *International Journal of Imaging Systems and Technology*, 11(2):101–116, 2000. doi: 10.1002/1098-1098(2000)11:2⟨101::aid-ima1⟩3.0.co;2-j. URL https://doi.org/10.1002/1098-1098(2000)11:2<101::aid-ima1>3.0.co;2-j.

Juergen Schmidhuber. Annotated History of Modern AI and Deep Learning, 2022. URL https://arxiv.org/abs/2212.11279.

Juergen Schmidhuber. Connectionists: Annotated History of Modern AI and Deep Learning. https://mailman.srv.cs.cmu.edu/pipermail/connectionists/2023-January/039120.html, 2023. Accessed on 09.03.2023.

Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, January 2015. doi: 10.1016/j.neunet.2014.09.003. URL https://doi.org/10.1016/j.neunet.2014.09.003.

Jürgen Schmidhuber. Critique of Honda Prize for Dr. Hinton. https://people.idsia.ch/~juergen/critique-honda-prize-hinton.html, 2020. Accessed on 09.03.2023.

Terrence J Sejnowski. *The Deep Learning Revolution*. The MIT Press. MIT Press, London, England, September 2018.

Terrence J. Sejnowski. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 117(48):30033–30038, January 2020. doi: 10.1073/pnas.1907373117. URL https://doi.org/10.1073/pnas.1907373117.

Terry Sejnowski. Connectionists: Scientific Integrity, the 2021 Turing Lecture, etc. https://mailman.srv.cs.cmu.edu/pipermail/connectionists/2022-January/037101.html, 2022. Accessed on 07.05.2023.

SIGGRAPH. History Community - ACM SIGGRAPH. https://www.siggraph.org/siggraph-365/history-community/, Accessed on 07.05.2024a.

SIGGRAPH. History — SIGGRAPH 2023. https://s2023.siggraph.org/program/history/, Accessed on 07.05.2024b.

Initial Signatories. An Open Letter to the Communications of the ACM. https://researchers.one/articles/20.12.00004v2, 2020. Accessed on 09.03.2023.

Shagun Sodhani, Mayoore S Jaiswal, Lauren Baker, Koustuv Sinha, Carl Shneider, Peter Henderson, Joel Lehman, and Ryan Lowe. Ideas for Improving the Field of Machine Learning: Summarizing Discussion from the NeurIPS 2019 Retrospectives Workshop. *arXiv preprint arXiv:2007.10546*, 2020. URL https://arxiv.org/abs/2007.10546.

Taylor Soper. Retired UW computer science professor embroiled in Twitter spat over AI ethics and 'cancel culture'. https://www.geekwire.com/2020/retired-uw-computer-science-professor-embroiled-twitter-spat-ai-ethics-cancel-culture/, 2020. Accessed on 09.03.2023.

Ivana Stanko. The Architectures of Geoffrey Hinton. In Sandro Skansi, editor, *Guide to Deep Learning Basics: Logical, Historical and Philosophical Perspectives*, pages 79–92. Springer International Publishing, Cham, 2020. ISBN 978-3-030-37591-1. doi: 10.1007/978-3-030-37591-1_8. URL https://doi.org/10.1007/978-3-030-37591-1_8.

James Steinhoff. Industrializing intelligence: A political economic history of the ai industry. In *Automation and Autonomy: Labour, Capital and Machines in the Artificial Intelligence Industry*, pages 99–131. Springer International Publishing, Cham, 2021. ISBN 978-3-030-71689-9. doi: 10.1007/978-3-030-71689-9_4. URL https://doi.org/10.1007/978-3-030-71689-9_4.

Charles C. Tappert. Who Is the Father of Deep Learning? In *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 343–348. IEEE, December 2019. doi: 10.1109/csci49370.2019.00067. URL https://doi.org/10.1109/csci49370.2019.00067.

Google Trends. "machine learning" - Google Trends. https://trends.google.com/trends/explore?date=all&q=%22machine%20learning%22&hl=en, 2023. Accessed on 03.03.2025.