Rethinking supervised learning: insights from biological learning and from calling it by its name

Anonymous Author(s)
Affiliation
Address
email

Abstract

The renaissance of artificial neural networks was catalysed by the success of classification models, tagged by the community with the broader term *supervised learning*. The extraordinary results gave rise to a hype loaded with ambitious promises and overstatements. Soon the community realised that the success owed much to the availability of thousands of labelled examples. And supervised learning went from glory to shame. Some criticised deep learning as a whole and others proclaimed that the way forward had to be "alternatives" to supervised learning: *predictive learning*, *unsupervised learning*, *semi-supervised learning* and, more recently, *self-supervised learning*. These seem all brand names, rather than actual categories of a theoretically grounded taxonomy. Moreover, the call to banish supervised learning was motivated by the overstatement that humans learn with little or no supervision. Here, we review insights about learning and supervision in nature, revisit the notion that learning is *not* possible without supervision and discuss why we will make better progress if we just call it by its name.

1 Introduction

The re-emergence of deep learning during the last decade due to the noteworthy achievements of artificial neural networks (ANN) built up a sort of philosophy that nearly anything could be automatically learnt from data without human intervention, in contrast to the previous approaches: "[hand designing good feature extractors, engineering skill and domain expertise] can all be avoided if good features can be learned automatically using a general-purpose learning procedure. This is the key advantage of deep learning" (LeCun et al., 2015).

However, this claim was overstated and can be misleading, since it misses the fact that the success of deep learning has required iterative hand design of network architectures, that demanded high and collective engineering skill and large doses of interdisciplinary domain expertise. Additionally, deep learning owes much to the immense computational power poured into training artificial networks (Amodei & Hernandez, 2018; Schwartz et al., 2019) and to the human effort of manually labelling thousands of images and other data modalities (Russakovsky et al., 2015; Cao et al., 2018).

At the same time, the need for much data has been used for motivating new research avenues and goals of deep learning that can be equally misleading and overly ambitious: "learning a class from a single labelled example", based on the statement that "humans learn new concepts with very little supervision", but "the standard supervised deep learning paradigm does not offer a satisfactory solution for learning new concepts rapidly from little data" (Vinyals et al., 2016). Some of the work that followed up these proposals certainly brought about promising methods and insights, but at the price of a publication landscape full of confusing terminology and research directions.

Submitted to 2nd Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM).

Here, we reflect upon fundamental concepts in machine learning such as supervision and inductive biases, which in spite of resting on solid theoretical grounds, are at the core of misconceptions and overstatements about deep learning commonly seen in the literature. We revisit insights from biological learning and from statistical learning theory with the aim of tempering certain claims and promises of deep learning and mitigating the confusion over the terminology and research directions.

2 Rethinking supervised learning

The realisation that the success of deep learning was largely due to the availability of huge labelled data sets prompted various reactions: some strongly questioned the usefulness of the algorithms (Marcus, 2018); some delved into the question of whether neural networks generalise beyond or simply memorise the training examples (Zhang et al., 2017; Arpit et al., 2017); and some opened new research horizons arguing that models should generalise from a few examples (Vinyals et al., 2016). As a consequence, multiple research programmes, with various brands, followed up with the aim of minimising or removing the need for "supervision" to train neural networks: few-shot, one-shot, zero-shot learning; predictive, unsupervised, semi-supervised, self-supervised learning and so on.

Improving the efficiency of learning algorithms should be a priority and related approaches had been subject of machine learning research before the explosion of deep learning. However, the current publication and discussion trends in the field denote overambitious promises, based partly on misconceptions and overstatements about learning in nature, and amplified by overselling nomenclature. While much of the research output does provide us with useful techniques and insight, the confusion amidst the hype run the risk of leading many astray.

In this section we aim to reassess these promises by first drawing insights from biological learning in the hope for setting upper bounds—at least in the near future—to our expectations about deep learning; and subsequently rethink the terminology around supervised learning commonly employed in today's deep learning literature.

2.1 Supervised biological learning

The link between artificial intelligence—specifically artificial neural networks (Rosenblatt, 1958; Fukushima & Miyake, 1982)—and biological learning systems is intrinsic to the field, as one long-term goal of artificial intelligence is to mirror the capabilities of human intelligence. However, these capabilities are sometimes overestimated. One example is the argument that intelligence in nature evolves without supervision and that humans and other animals learn—to visually categorise objects, for example—with little or no supervision (Vinyals et al., 2016; Marcus, 2018; Morgenstern et al., 2019). In what follows, we will discuss three aspects of biological learning to argue against this view, in order to gain insights that better inform our progress in machine learning: first, we will discuss how generalisation requires exposure to relevant *training data*; second, the role of evolution and brain development; third, the variety of supervised signals that the brain has access to.

Generalisation requires exposure to relevant training data First, in the argument that machine learning models should generalise from a few examples, there seems to be a promise or aspiration that future better methods will be able to perform robust visual object categorisation among many object classes after being trained on one or a few examples per class. While a primary objective is to develop techniques that efficiently extract the maximum possible information from the available examples, we should also remind ourselves that no machine learning algorithm can robustly learn anything that cannot be inferred from the data it has been trained on. Although this may seem to contradict certain current trends and statements in the literature, we should also bear in mind that learning in nature is not too different.

Even though the human visual system, for example, is remarkably robust, its capabilities are optimised for the tasks it needs to perform and largely shaped by experience, that is the *training data*—and years of evolution, as we will discuss below. For instance, from the literature on human visual perception, we know that object recognition is viewpoint dependent (Tarr et al., 1998). A well-studied property of human vision is that our face recognition ability is severely impaired if faces are presented upside down (Yin, 1969; Valentine, 1988). Setting aside the specific complexity of face processing in the brain, a compelling explanation for this impairment is that we are simply not used to seeing and recognising inverted faces. More generally, human perception of objects and our recognition ability

is greatly affected when we see objects from unfamiliar viewpoints (Edelman & Bülthoff, 1992; Bülthoff & Newell, 2006; Milivojevic, 2012).

Furthermore, although better than the *one-shot* or *few-shot* generalisation of current ANNs, humans also have limited ability to recognise novel classes (Morgenstern et al., 2019). Interestingly, experiments with certain novel classes of objects known as Greebles showed that with sufficient training, humans can acquire expertise in recognising new objects from different viewpoints, even making use of an area of the brain—the fusiform face area—that typically responds strongly with face stimuli (Gauthier et al., 1999). This provides evidence that generalisation to multiple viewpoints is possible but only developed after exposure to similar conditions.

The role of evolution and brain development Second, the commonplace comparison of artificial neural networks with the brain often misses a fundamental component of biology, recently brought to the fore by Zador (2019) and Hasson et al. (2020), although considered since the early days of artificial intelligence (Turing, 1968): the role that millions of years of evolution have played in developing the nervous systems of organisms in nature, including the human brain.

The properties and capabilities of ANNs trained for visual object categorisation are often contrasted with those of the adult brain (Yamins et al., 2014; Khaligh-Razavi & Kriegeskorte, 2014; Güçlü & van Gerven, 2015; Rajalingham et al., 2018; Geirhos et al., 2018). However, while artificial neural networks are typically trained from scratch[1], from random initialisation, a large part of the brain connectivity is encoded genetically and certain properties of the visual cortex are known to be innate, that is developed without prior exposure to visual stimuli (Zador, 2019).

These comparisons, although interesting, implicitly assume that the process by which a neural network learns its first visual categories from random initialisation is equivalent to the development of a brain from birth to maturity. It seems reasonable to at least consider that training a neural network on simple tasks may be more similar to learning in the infant brain (Harwerth et al., 1986; Atkinson, 2002; Gelman & Meyer, 2011), and probably even more to the process of evolution. This may be another reason to temper the expectations that neural networks be able to generalise robustly from a few examples, without *hard-wiring* some of the innate properties of the brain (Lindsey et al., 2019; Malhotra et al., 2020) or, alternatively, simulating part of the evolutionary process.

Supervised signals in the brain Third, another commonly found argument has it that children, and other animals in general, learn robust object recognition without supervision: "a child can generalize the concept of 'giraffe' from a single picture in a book" (Vinyals et al., 2016). First of all, we should recall again the role of evolution, which can be interpreted as a pre-trained model, optimised through millions of years of data with natural selection as a supervised signal (Zador, 2019). Second, we will argue against the very claim that children learn in fully unsupervised fashion.

Obviously, the kind of supervision that humans make use of is not that of classification algorithms: we do not see a class label on top of every object we look at. However, we receive supervision from multiple sources. Even though not for every visual stimulus, children do frequently receive information about the object classes they see—for instance, parents would point at objects and name them, then we learn how to read, and so on. Furthermore, humans usually follow a guided hierarchical learning: children do not directly learn to tell apart breeds of dogs, but rather start with umbrella terms and then progressively learn down the class hierarchy (Bornstein & Arterberry, 2010). Hasson et al. (2020) mention other examples of supervision from *social cues*, that is from other humans, such as learning to recognise individual faces, produce grammatical sentences, read and write; as well as from embodiment and action, such as learning to balance the body while walking or grasping objects. In all these actions, we can identify a supervised signal that surely influences learning in the brain (Shapiro, 2012).

Moreover, besides this kind of explicit supervision, the brain certainly makes use of more subtle, implicit supervised signals, such as temporal stability (Becker, 1999; Wyss et al., 2003): The light that enters the retina is not a random signal from a sequence of rapidly changing arbitrary photos, but a

[1] Some interesting and promising areas in machine learning research deviate from this standard approach. For example, transfer learning and domain adaptation study the potential of features learnt on one task to be reused in different, related tasks (Zhuang et al., 2019), and continual learning studies the ways in which machine learning models can indefinitely sustain the acquisition of new knowledge without detriment of the previously learnt tasks (Mundt et al., 2020). These approaches are inspired by biological learning or share interesting properties with it.

highly coherent and regular flow of slowly changing stimuli, especially at the higher, semantical level (Kording et al., 2004). At the very least, this is how we perceive it and if such a smooth perception is a consequence rather than a cause, then it should be, again, a by-product of a long process of evolution.

In the light of these insights from evolutionary theory and the explicit and implicit supervision that drives biological learning, we argue that we should temper the claims and aspirations that artificial neural networks should learn without supervision and from very few examples. Instead, we may benefit from rethinking the concept of supervision, embrace it and try to incorporate the forms of supervision present in nature into machine learning algorithms.

2.2 Supervised machine learning

If we open a machine learning textbook (Murphy, 2012; Abu-Mostafa et al., 2012; Goodfellow et al., 2016), we will most surely find a taxonomy of learning algorithms with a clear distinction between *supervised* and *unsupervised* learning. However, while this separation can be useful, the boundaries are certainly not clear. As a matter of fact, strictly speaking, unsupervised learning is an illusion. If we take a look at the deep learning literature of the past years, we will also find abundant work on some variants supposedly *in between*—semi-supervised learning, self-supervised learning, etc.—whose definitions are all but clear.

Catastrophic forgetting of old concepts If we recall a classical result in statistical learning theory and inference, the *no free lunch* theorem (Wolpert, 1996), no learning algorithm is better than any other at classifying unobserved data points, when averaged over all possible data distributions. Therefore, we need to constrain the distributions or, in other words, introduce prior knowledge—that is *supervision*. Recently, Locatello et al. (2018) obtained a related result for the case of unsupervised learning of disentangled representations: without inductive biases for both the models and the data sets, unsupervised disentanglement learning is impossible. These results are purely theoretical and have limited impact on real world applications Giraud-Carrier & Provost (2005), precisely because in practice we use multiple inductive biases, even when we do so-called unsupervised learning.

In a strict sense, even the classical, *purely* unsupervised methods, such as independent component analysis or nearest neighbours classifiers, make use of inductive biases, such as independence or minimum distance, respectively. Without inductive bias, learning is not possible. Although this is not news, the terminology used in the machine learning literature seems to neglect these nuances and evidences that the field suffers catastrophic forgetting of well-established notions.

The brands of *alt-supervised* learning Particularly in deep learning and computer vision, the term *supervised* learning is often used to actually refer to *classification*, that is models trained on examples labelled according to, for instance, the object classes. In turn, *unsupervised* learning is used for any model that does not use the labels, regardless of what other kind of supervision it may use. Further, the term *semi-supervised* learning refers in practice to models that are trained with a fraction of the labels, but are tested on the same classification benchmarks. More recently, the term *self-supervised* learning has gained much popularity[2], referring to models that are trained on tasks other than the standard task defined by classification labels.

Supervision comes in different flavours Some of the methods proposed under these brand names are certainly useful—that is not the subject of criticism of this work—but the terminology is overly confusing and unnecessary. First, there are not consensus definitions of these categories. A newcomer would easily fall into a scientific rabbit hole trying to tell these terms apart through publications—not to mention if they incorporated social media discussions into their endeavour. Second, from a theoretical point of view, both the conventional classification models and the recent wave of "self-supervised" tasks can all be formalised within the category of supervised learning. Only that supervision comes in different flavours, not only as classification labels.

We have seen examples of different forms of supervision used by humans and other animals. In machine learning, the field focused for many years on a few loss functions, such as classification and

[2]Especially, after Yann LeCun's keynote presentation at ISSCC in February 2019. In December 2016, he had titled his NeurIPS keynote presentation "Predictive Learning", to refer to similar ideas, pointing out that "that's what many people mean by unsupervised learning". Later, he wrote: "I Now call it 'self-supervised learning', because 'unsupervised' is both a loaded and confusing term".

simple forms of regression. The relatively recent explosion of deep learning has brought about the development of several libraries for automatic differentiation (Baydin et al., 2017), which in turn have enabled the proposal of multiple loss functions with various types of supervision (see (Jing & Tian, 2020) for a recent review) that can easily be optimised numerically by stochastic gradient descent and artificial neural networks. This opens in fact a great opportunity to incorporate new inductive biases from biological learning into machine learning algorithms, beyond classification labels. However, we contend that this is also supervised learning and we should call it by its name.

The terms *self*- and *semi*-supervised imply that the methods use *less* supervision, but, in fact, they should have been called, if anything, *hyper-supervised*[3] learning. Pure classifiers use the labels as main supervised signal, but self- and semi-supervised methods often make use of a wide range of surrogate tasks with supervised signals defined by humans. A well-known example is image data augmentation: Although at first glance it may seem a naïve technique, it actually encodes rich prior knowledge about visual perception and that is why it outperforms explicit regularisation methods with more naïve inductive biases Hernández-García & König (2018) and is widely used in "semi-supervised" tasks Laine & Aila (2016).

3 Conclusion

In this paper, we have first discussed some of the overambitious promises and misconceptions of the deep learning hype around the need for human intervention and supervision in machine learning. In particular, we have critically reviewed the various terms that are currently used to refer to supposed alternatives to supervised learning: semi-, self- and unsupervised learning, among others. In sum, we pointed out that all these approaches are in fact supervised learning—though not necessarily classification—and the machine learning (research) community would benefit from using more rigorous, less overselling nomenclature.

Supervision is not evil. It is at the core of statistical learning theory: learning is impossible without inductive biases or supervision. But it has different flavours, not only classification. Neither is deep learning some sort of exceptional solution to learn without human intervention and supervision, nor is it a hopeless model class because it requires large data sets (Marcus, 2018). The human visual system is exposed to a lot of stimuli too. One exceptional advantage of deep learning is precisely that it is possible to effectively optimise different learning objectives, almost end-to-end, from large collections of nearly naturalistic sensory signals, such as digital images (Saxe et al., 2020). While other models are known to scale poorly as the amount of data increases, neural networks excel at fitting the training data and interpolating on unseen examples (Belkin et al., 2019; Hasson et al., 2020). This is a feature, not a bug. But we will make better progress if we exploit these advantages of deep learning without neglecting that supervision will always be necessary—the critical goal is how to best incorporate it and exploit it.

We have drawn parallels from cognitive science and human perception to contend that learning in nature also requires supervision, where it too has multiple forms. Even evolution can be regarded as an optimisation process where natural selection is the supervised signal (Zador, 2019). We have argued, as others have before, that these insights from biology and neuroscience offer a great opportunity for machine learning research to draw inspiration and calibrate its compass.

Finally, we recall that most of the learning theory has been developed for simple loss functions such as binary classification or mean squared error regression, but certain methods successfully used in practice today escape the available theory. Given the success of this kind of more complex supervised objectives and their connection with perception, the study of these methods from a theoretical point of view might be a fruitful direction for future work.

Broader Impact

Since this article does not present a new method or results from data sets, potential risks of "bias in the data" or "failure of the system" do not apply. As a critical review of current trends in the field, some researchers could potentially be affected depending on their work. We declare that we do not

[3]The authors explicitly discourage the addition of a new term to the already too confusing list.

intend to negatively affect any individual researcher, but rather potentially improve scientific value through a constructive reflection.

References

Abu-Mostafa, Y. S., Magdon-Ismail, M., and Lin, H.-T. *Learning from data*. AMLBooks, 2012.

Amodei, D. and Hernandez, D. AI and compute. Accessed: 2020-09-28, 2018. URL `https://openai.com/blog/ai-and-compute/`.

Arpit, D. et al. A closer look at memorization in deep networks. *arXiv preprint arXiv:1706.05394*, 2017.

Atkinson, J. *The developing visual brain*. Oxford Scholarship Online, 2002.

Baydin, A. G., Pearlmutter, B. A., Radul, A. A., and Siskind, J. M. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research (JMLR)*, 2017.

Becker, S. Implicit learning in 3D object recognition: The importance of temporal context. *Neural Computation*, 1999.

Belkin, M., Hsu, D., Ma, S., and Mandal, S. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences (PNAS)*, 2019.

Bornstein, M. H. and Arterberry, M. E. The development of object categorization in young children: Hierarchical inclusiveness, age, perceptual attribute, and group versus individual analyses. *Developmental Psychology*, 2010.

Bülthoff, I. and Newell, F. N. The role of familiarity in the recognition of static and dynamic objects. *Progress in Brain Research*, 2006.

Cao, Q., Shen, L., Xie, W., Parkhi, O. M., and Zisserman, A. VGGFace2: A dataset for recognising faces across pose and age. In *IEEE International Conference on Automatic Face & Gesture Recognition*. 2018.

Edelman, S. and Bülthoff, H. H. Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, 1992.

Fukushima, K. and Miyake, S. Neocognitron: A new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognition*, 1982.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., and Gore, J. C. Activation of the middle fusiform'face area'increases with expertise in recognizing novel objects. *Nature Neuroscience*, 1999.

Geirhos, R., Temme, C. R., Rauber, J., Schütt, H. H., Bethge, M., and Wichmann, F. A. Generalisation in humans and deep neural networks. In *Advances in neural information processing systems*, 2018.

Gelman, S. A. and Meyer, M. Child categorization. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2011.

Giraud-Carrier, C. and Provost, F. Toward a justification of meta-learning: Is the no free lunch theorem a show-stopper. In *Proceedings of the ICML-2005 Workshop on Meta-learning*, 2005.

Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. MIT press, 2016.

Güçlü, U. and van Gerven, M. A. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, 2015.

Harwerth, R. S., Smith, E. L., Duncan, G. C., Crawford, M., and Von Noorden, G. K. Multiple sensitive periods in the development of the primate visual system. *Science*, 1986.

Hasson, U., Nastase, S. A., and Goldstein, A. Direct fit to nature: An evolutionary perspective on biological and artificial neural networks. *Neuron*, 2020.

Hernández-García, A. and König, P. Data augmentation instead of explicit regularization. *arXiv preprint arXiv:1806.03852*, 2018.

Jing, L. and Tian, Y. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.

Khaligh-Razavi, S.-M. and Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLOS Computational Biology*, 2014.

Kording, K. P., Kayser, C., Einhauser, W., and Konig, P. How are complex cell properties adapted to the statistics of natural stimuli? *Journal of Neurophysiology*, 2004.

Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.

LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*, 2015.

Lindsey, J., Ocko, S. A., Ganguli, S., and Deny, S. A unified theory of early visual representations from retina to cortex through anatomically constrained deep CNNs. In *International Conference on Learning Representations (ICLR), arXiv:1901.00945*, 2019.

Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. *arXiv preprint arXiv:1811.12359*, 2018.

Malhotra, G., Evans, B. D., and Bowers, J. S. Hiding a plane with a pixel: examining shape-bias in CNNs and the benefit of building in biological constraints. *Vision Research*, 2020.

Marcus, G. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018.

Milivojevic, B. Object recognition can be viewpoint dependent or invariant–it's just a matter of time and task. *Frontiers in Computational Neuroscience*, 2012.

Morgenstern, Y., Schmidt, F., and Fleming, R. W. One-shot categorization of novel object classes in humans. *Vision Research*, 2019.

Mundt, M., Hong, Y. W., Pliushch, I., and Ramesh, V. A wholistic view of continual learning with deep neural networks: Forgotten lessons and the bridge to active and open world learning. *arXiv preprint arXiv:2009.01797*, 2020.

Murphy, K. P. *Machine learning: a probabilistic perspective*. MIT Press, 2012.

Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., and DiCarlo, J. J. Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience*, 2018.

Rosenblatt, F. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 1958.

Russakovsky, O. et al. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 2015.

Saxe, A., Nelli, S., and Summerfield, C. If deep learning is the answer, then what is the question? *arXiv preprint arXiv:2004.07580*, 2020.

Schwartz, R., Dodge, J., Smith, N. A., and Etzioni, O. Green ai. *arXiv preprint arXiv:1907.10597*, 2019.

Shapiro, L. *Embodied cognition*. Oxford Handbooks Online, 2012.

Tarr, M. J., Williams, P., Hayward, W. G., and Gauthier, I. Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience*, 1998.

Turing, A. M. *Cybernetics; (Key papers)*. University Park Press, 1968.

Valentine, T. Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, 1988.

Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

Wolpert, D. H. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 1996.

Wyss, R., König, P., and Verschure, P. F. Invariant representations of visual patterns in a temporal population code. *Proceedings of the National Academy of Sciences (PNAS)*, 2003.

Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences (PNAS)*, 2014.

Yin, R. K. Looking at upside-down faces. *Journal of Experimental Psychology*, 1969.

Zador, A. M. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 2019.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations (ICLR), arXiv:1611.03530*, 2017.

Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. A comprehensive survey on transfer learning. *arXiv preprint arXiv:1911.02685*, 2019.