

LEVELS OF ANALYSIS FOR MACHINE LEARNING

Jessica B. Hamrick
DeepMind
jhamrick@google.com

Shakir Mohamed
DeepMind
shakir@google.com

ABSTRACT

Machine learning is currently involved in some of the most vigorous debates it has ever seen. Such debates often seem to go around in circles, reaching no conclusion or resolution. This is perhaps unsurprising given that researchers in machine learning come to these discussions with very different frames of reference, making it challenging for them to align perspectives and find common ground. As a remedy for this dilemma, we advocate for the adoption of a common conceptual framework which can be used to understand, analyze, and discuss research. We present one such framework which is popular in cognitive science and neuroscience and which we believe has great utility in machine learning as well: *Marr’s levels of analysis*. Through a series of case studies, we demonstrate how the levels facilitate an understanding and dissection of several methods from machine learning. By adopting the levels of analysis in one’s own work, we argue that researchers can be better equipped to engage in the debates necessary to drive forward progress in our field.

1 CONCEPTUAL FRAMEWORKS

The last few months and years have seen researchers embroiled in heated debates touching on fundamental questions in machine learning: what, exactly, is “deep learning”? Can it ever be considered “symbolic”? Is it sufficient to achieve artificial general intelligence? Can it be trusted in high-stakes, real-world applications? These are important questions, yet the discussions surrounding them seem unable to move forward productively. This is due, in part, to the lack of a common framework within which argument and evidence are grounded and compared. Early researchers in computer science, engineering and cognitive science struggled in this same way, and were led to develop several *conceptual frameworks*: mental models that align researchers’ understanding of how their work and philosophies fit into the larger goals of their field and science more broadly (e.g. Chomsky, 1965; Marr & Poggio, 1976; Newell & Simon, 1976; Marr, 1982; Newell, 1981; Arbib, 1987; Anderson, 1990). Our proposition is to reinvigorate this former tradition of debate by making the use of conceptual frameworks a core part of machine learning research.

We focus in this paper on the use of *Marr’s levels of analysis*, a conceptual framework popular in cognitive science and neuroscience. Marr (1982) identified a three-layer hierarchy for describing and analyzing a computational system:

Computational level. *What is the goal of a system, what are its inputs and outputs, and what mathematical language can be used to express that goal?* As an example, let us consider the case of natural language. In one view, Chomsky (1965) famously argued that the purpose of language is to structure thought by expressing knowledge in a compositional and recursive way.

Algorithmic or representational level. *What representations and algorithms are used to achieve the computational-level goal?* Under Chomsky’s view of language, we might hypothesize that a language processing system uses symbolic representations like syntax trees and that those representations are formed using some form of linguistic parsing.

Implementation level. *How is the system implemented, either physically or in software?* For example, we might ask how a neural circuit in the brain could represent syntax trees using population-based codes, or how a parser should be implemented efficiently in code.

Each level of analysis allows us to identify various hypotheses and constraints about how the system operates, and to highlight areas of disagreement. For example, others have argued that the computational-level goal of language is not to structure thought, but to communicate with others (Tomasello, 2010). At the algorithmic level, we might argue that syntax trees are an impoverished

representation for language, and that we must also consider dependency structures, information compression schemes, context, pragmatics, and so on. And, at the implementation level, we might consider how things like working memory or attention constrain what algorithms can be realistically supported.

2 CASE STUDIES IN USING LEVELS OF ANALYSIS

We use four case studies to show how Marr’s levels of analysis can be applied to understanding methods in machine learning, choosing examples that span the breadth of research in the field. While the levels have traditionally been used to reverse-engineer biological behavior, here we use them as a conceptual device that facilitates the structuring and comparison of machine learning methods. Additionally, we note that our classification of various methods at different levels of analysis may not be the only possibility, and readers may find themselves disagreeing with us. This disagreement is a desired outcome: the act of applying the levels forces differing assumptions out into the open so they can be more readily discussed.

2.1 DEEP Q-NETWORKS (DQN)

DQN (Mnih et al., 2015) is a deep reinforcement learning algorithm originally used to train agents to play Atari games. At the **computational** level, the goal of DQN is to efficiently produce actions which maximize scalar rewards given observations from the environment. We can use the language of reinforcement learning—and in particular, the Bellman equation—to express this goal (Sutton & Barto, 2018). At the **algorithmic** level, DQN performs off-policy one-step temporal difference (TD) learning (i.e., Q-learning), which can be contrasted with other choices such as SARSA, n -step Q-learning, $TD(\lambda)$, and so on. All of these different algorithmic choices are different realizations of the same goal: to solve the Bellman equation. At the **implementation** level, we consider how Q-learning should actually be implemented in software. For example, we can consider the neural network architecture, optimizer, and hyperparameters, as well as components such as a target network and replay buffer. We could also consider how to implement Q-learning in a distributed manner (Horgan et al., 2018).

The levels of analysis when applied to machine learning systems help us to more readily identify assumptions that are being made and to critique those assumptions. For example, perhaps we wish to critique DQN for failing to *really* know what objects are, such as what a “paddle” or “ball” is in Breakout (Marcus, 2018). While on the surface this comes across as a critique of deep reinforcement learning more generally, by examining the levels of analysis, it becomes clear that this critique could actually be made in very different ways at two levels of analysis:

1. At the computational level, the goal of DQN is not to learn about objects but to maximize scalar reward in a game. Yet, if we care about the system understanding objects, then perhaps the goal should be formulated such that discovering objects is part of the objective itself (e.g. Burgess et al., 2019; Greff et al., 2019). Importantly, the fact that the agent uses deep learning is orthogonal to the question of what the goal is.
2. At the implementation level, we might interpret the critique as being about the method used to implement the Q-function. Here, it becomes relevant to talk about deep learning. DQN uses a convolutional neural network to convert visual observations into distributed representations, but we could argue that these are not the most appropriate for capturing the discrete and compositional nature of objects (e.g. Battaglia et al., 2018; van Steenkiste et al., 2019).

Depending on whether the critique is interpreted at the computational level or implementation level, one might receive very different responses, thus leading to a breakdown of communication. However, if the levels were to be used and explicitly referred to when making an argument, there would be far less room for confusion or misinterpretation and far more room for productive discourse.

2.2 CONVOLUTIONAL NEURAL NETWORKS

One of the benefits of the levels of analysis is that they can be flexibly applied to any computational system, even when that particular system might itself be a component of a larger system. For example, an important component of the DQN agent at the implementation level is a visual frontend which processes observations and maps them to Q-values (Mnih et al., 2015). Just as we can analyze the whole DQN agent using the levels of analysis, we can do the same for DQN’s visual system.

At the **computational** level, the goal of a vision module is to map spatial data (as opposed to, say, sequential data or graph-structured data) to a compact representation of objects and scenes that is invariant to certain types of transformations (such as translation). At the **algorithmic** level, feed-forward convolutional networks (Fukushima, 1988; LeCun et al., 1989) are one type of procedure that processes spatial data while maintaining translational invariance. Within the class of CNNs, there are many different versions, such as AlexNet (Krizhevsky et al., 2012), ResNet (He et al., 2016), or dilated convolutions (Yu & Koltun, 2015). However, this need not be the only choice. For example, we could consider a relational (Wang et al., 2018) or autoregressive (Oord et al., 2016) architecture instead. Finally, at the **implementation** level, we can ask how to efficiently implement the convolution operation in hardware. We could choose to implement our CNN on a CPU, a single GPU, or multiple GPUs. We may also be concerned with questions about what to do if the parameters or gradients of our network are too large to fit into GPU memory, how to scale to much larger batch sizes, and how to reduce latency.

Analyzing the CNN on its own highlights how the levels can be applied flexibly to many types of methods: we can “zoom” our analysis in and out to focus understanding and discussion of different aspects of larger systems as needed. This ability is particularly useful for analyzing whether a component of a larger system is really the right one by comparing the role of a component in a larger system to its computational-level goal in isolation. For example, as we concluded above, the computational-level goal of a CNN is to process spatial data while maintaining translational invariance. This might be appropriate for certain types of goals (e.g., object classification) but not for others in which translational invariance is inappropriate (e.g., object localization).

2.3 SYMBOLIC REASONING ON GRAPHS

A major topic of debate has been the relationship between symbolic reasoning systems and distributed (deep) learning systems. The levels of analysis provide an ideal way to better illuminate the form of this relationship. As an illustrative example, let us consider the problem of solving an NP-complete problem like the Traveling Salesman Problem (TSP), a task that has traditionally been approached symbolically. At the **computational** level, given a complete weighted graph (i.e. fully-connected edges with weights), the goal is to find the minimum-weight path through the graph that visits each node exactly once. Although finding an exact solution to the TSP takes (in the worst case) exponential time, we could formulate part of the goal as finding an efficient solution which returns near-optimal or approximate solutions. At the **algorithmic** level, there are many ways we could try to both represent and solve the TSP. As it is an NP-complete problem, we could choose to transform it into other types of NP-complete problems (such as graph coloring, the knapsack problem, or boolean satisfiability). A variety of algorithms exist for solving these different types of problems, with a common approach being heuristic search. While heuristics are typically handcrafted at the **implementation** level using a symbolic programming language, they could also be implemented using deep learning components (e.g. Vinyals et al., 2015; Dai et al., 2017).

Performing this analysis shows how machine learning systems may simultaneously implement computations with *both* symbolic and distributed representations at different levels of analysis. Perhaps, then, discussions about “symbolic reasoning” versus “deep learning” or (“hybrid systems” versus “non-hybrid systems”) are not the right focus because *both* symbolic reasoning and deep learning already coexist in many architectures. Instead, we can use the levels to productively steer discussions to where they can have more of an impact. For example, we could ask: is the logic of the algorithmic level the right one to achieve the computational-level goal? Is that goal the one we care about solving in the first place? And, is the architecture used at the implementation level appropriate for learning or implementing the desired algorithm?

2.4 MACHINE LEARNING IN HEALTHCARE

The last few case studies have focused on applying Marr’s levels to widely-known algorithms. However, an important component of any machine learning system is also the data on which it is trained and its domain of application. We turn to an application of machine learning in real-world settings where data plays a significant role. Consider a clinical support tool for managing electronic health records (EHR) (e.g. Wu et al., 2010; Tomašev et al., 2019). At the **computational** level, the goal of this system is to improve patient outcomes by predicting patient deterioration during the course of hospitalization. Performing this computational-level analysis encourages us to think deeply about all facets of the overall goal, such as: what counts as a meaningful prediction that gives sufficient time to act? Are there clinically-motivated labels available? How should the system deal with imbalances

in clinical treatments and patient populations? What unintended or undesirable consequences might arise from deployment of such a system? The **algorithmic** level asks what procedures should be used for collecting and representing the data in order to achieve the computational-level goal. For example, how are data values encoded (e.g., censoring extreme values, one-hot encodings, using derived summary-statistics)? Is the time of an event represented using regular intervals or not? And, what methods allow us to develop longitudinal models of such high-dimensional and highly-sparse data? This is where familiar models of gradient-boosted trees or logistic regression enter. Finally, at the **implementation** level, questions arise regarding how the system is deployed. For example, we may need to consider data sandboxing and security, inter-operability, the role of open-standards, data encryption, and interactions and behavioural changes when the system is deployed.

In addition to facilitating debate, applying Marr’s levels of analysis in a domain like healthcare emphasizes important aspects of design decisions that need to be made. For example, the representation we choose at the algorithmic level can lock in choices at the beginning of a research program that has important implications for lower levels of analysis. Choosing to align the times when data is collected into six hour buckets will affect underlying structure and predictions, with implications on how performance is analyzed and for later deployment. Similarly, “alerting fatigue” is a common concern of such methods, and can be debated at the computational level by asking questions related to the broader interactions of the system with existing clinical pathways, or at the implementation level through mechanisms for alert-filtering. When we consider questions of bias in the data sets used, especially in healthcare, the levels can help us identify whether such bias is a result of clinical constraints at the computational level, or whether it is an artefact of algorithmic-level choices.

3 BEYOND MARR

As our case studies have demonstrated, Marr’s levels of analysis are a powerful conceptual framework that can be used to define, understand, and discuss research in machine learning. The reason that the levels work so well is that they allow researchers to approach discussion using a common frame-of-reference that highlights different choices or assumptions, streamlining discussion while also encouraging skepticism. We suggest readers try this exercise themselves: apply the levels to a recent paper that you enjoyed (or disliked!) as well as your own research. We have found this approach to be a useful way to understand our own specific papers as well as the wider research within which they sit. While the process is sometimes harder than expected, the difficulties that arise are often exactly in the places we have overlooked and where deeper investigation is needed.

Machine learning research often implicitly relies on several similar conceptual frameworks, although they are not often formalized in the way Marr’s levels are. For example, most deep learning researchers will recognize an *architecture-loss* paradigm: a given problem is studied in terms of a computational graph that processes and transforms data into some output, and a method that computes prediction errors and propagates them to drive parameter and distribution updates. Similarly, many software engineers will recognize the *algorithm-software-hardware* distinction: a problem can be understood in terms of the algorithm to be computed, the software which computes that algorithm, and the hardware on which the software runs. While both of these frameworks are conceptually useful, neither fully captures the full set of abstractions ranging from the computational-level goal to the physical implementation.

Of course, Marr’s levels are not without limitations or alternatives. Even in the course of writing this paper, we considered whether it would make sense to split the algorithmic level into two, mirroring the original four-level formulation put forth by Marr & Poggio (1976). Researchers from cognitive science and neuroscience have also spent considerable effort discussing alternate formulations of the levels (e.g. Sun, 2008; Griffiths et al., 2012; Poggio, 2012; Danks, 2013; Peebles & Cooper, 2015; Niv & Langdon, 2016; Yamins & DiCarlo, 2016; Krafft & Griffiths, 2018; Love, 2019). For example, Poggio (2012) argues that *learning* should be added as an additional level within Marr’s hierarchy. Additionally, Marr’s levels do not explicitly recognize the socially-situated role of computational systems, leading to proposals for alternate hierarchies that consider the interactions between computational systems themselves, as well as the socio-cultural processes in which those systems are embedded (Nissenbaum, 2001; Sun, 2008). The ideal conceptual framework for machine learning will need to go beyond Marr to also address such considerations. In the meantime, we believe that Marr’s levels are useful as a common conceptual lens for machine learning researchers to start with. Given such a lens, we believe that researchers will be better equipped to discuss and view their research, and to ultimately address the many deep questions of our field.

ACKNOWLEDGMENTS

We would like to thank M. Overlan, K. Stachenfeld, T. Pfaff, A. Santoro, D. Pfau, A. Sanchez-Gonzalez, M. Rosca, and two anonymous reviewers for helpful comments and feedback on this paper. Additionally, we would like to thank I. Momennejad, A. Erdem, A. Lampinen, F. Behbahani, R. Patel, P. Krafft, H. Ritz, T. Darmetko, and many others on Twitter for providing a variety of interesting perspectives and thoughts on this topic that helped to inspire this paper.

REFERENCES

- John R Anderson. *The adaptive character of thought*. Psychology Press, 1990.
- Michael A Arbib. Levels of modeling of mechanisms of visually guided behavior. *Behavioral and Brain Sciences*, 10(3):407–436, 1987.
- Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- Noam Chomsky. *Aspects of the Theory of Syntax*. MIT press, 1965.
- Hanjun Dai, Elias Khalil, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *Advances in Neural Information Processing Systems*, pp. 6348–6358, 2017.
- David Danks. Moving from levels & reduction to dimensions & constraints. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2013.
- Kunihiko Fukushima. Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural networks*, 1(2):119–130, 1988.
- Klaus Greff, Raphaël Lopez Kaufmann, Rishabh Kabra, Nick Watters, Chris Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. *arXiv preprint arXiv:1903.00450*, 2019.
- Thomas L Griffiths, Edward Vul, and Adam N Sanborn. Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4):263–268, 2012.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado Van Hasselt, and David Silver. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Peter Krafft and Tom Griffiths. Levels of analysis in computational social science. In *CogSci*, 2018.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- Bradley C Love. Levels of biological plausibility. *PsyArXiv preprint*, 2019. doi: <https://doi.org/10.31234/osf.io/49vea>.
- Gary Marcus. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*, 2018.

- David Marr. The philosophy and the approach. In *Vision: A computational investigation into the human representation and processing of visual information*. MIT Press, 1982.
- David Marr and Tomaso Poggio. From understanding computation to understanding neural circuitry, 1976. AI Memo 357.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Allen Newell. The knowledge level: presidential address. *AI magazine*, 2(2):1–1, 1981.
- Allen Newell and Herbert A Simon. Computer science as empirical inquiry: symbols and search. *Communications of the ACM*, 19(3):113–126, 1976.
- Helen Nissenbaum. How computer systems embody values. *Computer*, 34(3):120–119, 2001.
- Yael Niv and Angela Langdon. Reinforcement learning with marr. *Current opinion in behavioral sciences*, 11:67–73, 2016.
- Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.
- David Peebles and Richard P Cooper. Thirty years after marr’s vision: levels of analysis in cognitive science. *Topics in cognitive science*, 7(2):187–190, 2015.
- Tomaso Poggio. The levels of understanding framework, revised. *Perception*, 41(9):1017–1023, 2012.
- Ron Sun. Introduction to computational cognitive modeling. *Cambridge handbook of computational psychology*, pp. 3–19, 2008.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Michael Tomasello. *Origins of human communication*. MIT press, 2010.
- Nenad Tomašev, Xavier Glorot, Jack W Rae, Michal Zielinski, Harry Askham, Andre Saraiva, Anne Mottram, Clemens Meyer, Suman Ravuri, Ivan Protsyuk, et al. A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature*, 572(7767):116–119, 2019.
- Sjoerd van Steenkiste, Klaus Greff, and Jürgen Schmidhuber. A perspective on objects and systematic generalization in model-based rl. *arXiv preprint arXiv:1906.01035*, 2019.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in neural information processing systems*, pp. 2692–2700, 2015.
- Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7794–7803, 2018.
- Jionglin Wu, Jason Roy, and Walter F Stewart. Prediction modeling using ehr data: challenges, strategies, and a comparison of machine learning approaches. *Medical care*, pp. S106–S113, 2010.
- Daniel LK Yamins and James J DiCarlo. Eight open questions in the computational modeling of higher sensory cortex. *Current opinion in neurobiology*, 37:114–120, 2016.
- Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.