

ARTICLE

Social Learning in Neural Agent-Based Models

Igor Douven

Paris 1 Université Panthéon-Sorbonne, Paris, France

Email: igor.douven@gmail.com

(Received 08 February 2024; revised 12 June 2024; accepted 24 June 2024)

Abstract

Agent-based models (ABMs) are widely used to study how individual interactions shape collective behaviors. Critics argue that ABMs are often too simplistic to capture real-world complexities. We address this by integrating artificial neural networks into ABMs, focusing on enhancing the Hegselmann–Krause (HK) model. By using multilayer perceptrons as agents, we create more realistic ABMs that better reflect actual agents. This approach yields multiple models, as core elements of the HK model can be defined in various ways. We conduct two computational studies to compare these models with each other and with traditional individual-learning paradigms.

1 Introduction

Agent-based models (ABMs) have become a popular tool for studying macro-properties of social systems that, although typically arising from simple, micro-level interactions, cannot be fully understood by strictly analytical means. They are used across a range of domains, from economics and political science to epidemiology and urban planning (Crosscombe and Lawry 2016; Deffuant et al. 2000; Dittmer 2001; Douven and Hegselmann 2021; Lorig, Johansson, and Davidsson 2021; O'Connor and Weatherall 2019; Schelling 1971), and philosophers of science have recruited ABMs to argue that social learning is key to producing and acquiring scientific knowledge (Douven 2010; Glass and Glass 2021; Hegselmann et al. 2015; Huang, forthcoming; Kummerfeld and Zollman 2016; Olsson 2013; Olsson and Vallinder 2013; Rosenstock, O'Connor, and Bruner 2017; Zollman 2007, 2010).¹

Although popular, agent-based modeling has recently come under a cloud. According to various authors, ABMs tend to oversimplify agent behavior, decision-making processes, and environments, which—these authors argue—undermines

¹ The article has supplementary materials consisting of an online-only appendix and the data and code used for the simulations. The Jupyter notebook containing the code also includes extra analyses of the simulation outcomes and a short tutorial on defining neural networks using the Flux.jl package for the Julia language (Bezanson et al. 2017). All materials can be downloaded from the repository at <https://osf.io/fs29h/>.

their ability to adequately capture the complexity and variability of real-world behavior and thus to yield accurate predictions when applied to actual social processes (see, e.g., Borg et al. 2019; Cristelli 2014; Frey and Šešelja 2018, 2020; Rosenstock, O'Connor, and Bruner 2017; Šešelja 2019; Thicke 2020).

An obvious response to this critique is to make ABMs more realistic, which can be done, for instance, by letting interactions among agents be governed by more complex rules and by making the agents' environment more like the real world in relevant respects. Several agent-based COVID-19 models were successful because of this approach. Not only did these models capture relevant population differences (in terms of age, health status, social behavior, mobility patterns, etc.), as well as the resulting heterogeneity of the interactions among agents, but they were also able to incorporate data about the evolving pandemic almost in real time—features that made them valuable tools for policymakers in managing the pandemic (see, e.g., Adam 2020; Douven 2024). A different approach, taken in this article, is to make the agents *intrinsically* more humanlike by endowing them with some artificial form of intelligence. We aim to accomplish this by integrating ABMs with artificial neural networks (ANNs). The resulting neural agent-based models (NABMs) allow agents to learn from and adapt to their environment and interactions in a manner more akin to how humans learn and adapt.

This is still a very broad proposal, given the number of different ANN architectures on the market as well as the number of different ABMs with which they could be combined. As for ABMs, our focus will be on the Hegselmann–Krause (HK) model (Hegselmann and Krause 2002, 2006, 2015, 2019), a well-established framework for studying opinion dynamics that enjoys considerable popularity in philosophy and beyond. We will combine this model with multilayer perceptrons (MLPs), which are among the oldest types of ANN. The HK model captures the process of opinion formation and evolution within a society of agents, where the agents' opinions are influenced both by evidence obtained directly from the world and by the opinions of their peers, which are the agents whose opinions are close to the agents' own opinions. By populating the HK model with MLPs, we aim to simulate the behavior of agents that form their opinions, or more generally update their doxastic states, not on the basis of simple arithmetic operations (as the agents in the HK model do) but rather by leveraging an ability to process complex information in a humanlike way. Nevertheless, the dual updating mechanism characteristic of the HK model remains intact in our NABMs in that updating will still be one part data driven, one part based on social interactions, where the latter are achieved through either parameter averaging or prediction alignment, or both (in ways to be explained).

Our primary goal is to present what we believe to be a promising approach to making ABMs more realistic, thereby also addressing the recent critique of such models. A secondary goal is to assess what remains of the seeming support from ABMs for the efficacy of social learning when this issue is considered in more realistic settings. We provide some theoretical background on the two main components of our NABMs (the HK model and MLPs) in section 2. The HK-based NABMs are then presented in section 3. Sections 4 and 5 report computational studies conducted using these models, the first study centering on a classification task, the second involving probabilistic updating in the context of medical diagnostics. Both studies address the

question of the significance of social learning by comparing forms of such learning with each other and with individual learning.

As a preliminary note, we emphasize that the framework to be presented is meant as a blueprint for combining ABMs and ANNs generally and that our methodology can be adapted to integrate ABMs with ANNs more sophisticated and state of the art than MLPs, including the large language models (LLMs) that have been much in the limelight lately. At present, the requisite adaptations of the framework would encounter practical obstacles, for instance, due to the limited accessibility of cutting-edge LLMs—the most impressive ones being proprietary software—and the substantial computational resources required for training extensive numbers of larger networks. But anyone who has been following developments in the field of artificial intelligence will find it reasonable to expect that such challenges will be overcome sooner rather than later.

2 Theoretical background

2.1 The Hegselmann–Krause model

The HK model is among the most popular frameworks in the domain of agent-based computational modeling. While the model admits of a variety of interpretations (see, e.g., Hegselmann 2023), it is most commonly interpreted to encapsulate the interplay between two key aspects of human epistemic behavior: the assimilation of information from social peers and the direct acquisition of knowledge from empirical evidence. On this interpretation, it serves as a mathematical abstraction of opinion dynamics, where agents iteratively adjust their beliefs about the value of some parameter $\tau \in [0, 1]$, whose meaning remains unspecified.

Formally, each agent i starts at time 0 with an estimate $x_i(0)$ of τ and revises this estimate over discrete time steps, where at each time t the revision process is influenced by two primary factors: evidence about τ the agent receives directly from the world at t and the opinions of its peers at t , which are formally defined to be the agents within its bounded confidence interval (BCI) at t , that is, whose estimates of τ at t differ by no more than some small value ϵ from the agent's own estimate at t . Then agent i 's opinion concerning τ after the $(n + 1)$ st update (i.e., at time $n + 1$) is defined to be

$$x_i(n + 1) = \frac{1 - \alpha}{|X_i(n)|} \sum_{j \in X_i(n)} x_j(n) + \alpha\tau,$$

with $x_j(n)$ being the opinion of agent j after update n and

$$X_i(n) := \{j : |x_i(n) - x_j(n)| \leq \epsilon\}$$

the set of agents within agent i 's BCI after update n . The parameter $\alpha \in [0, 1]$ balances the weight given to social versus evidential information. (For illustrations, see the supplementary materials.)

A key virtue of the HK model is that it can be easily extended or adapted for the purpose of addressing specific research questions. For instance, researchers have explored scenarios with “noisy” evidence, where agents receive imperfect signals from the world (Douven 2010), and have considered agents with interval-valued beliefs to account for vagueness (Crosscombe and Lawry 2016) as well as agents that

can simultaneously hold beliefs about multiple issues (Jacobmeier 2004; Lorenz 2008; Pluchino, Latora, and Rapisarda 2006).

In line with the general critique of ABMs cited in the introduction, one could argue that the HK model, and even the aforementioned extensions of the model, features agents whose intellectual capacities are, for all we know, unrealistically impoverished. Proponents of the HK model could respond that this does not mean their model cannot be descriptively adequate at the macro-level, for instance, in predicting when a community of agents will reach a consensus and when it will not. While that is true, we believe that a more productive, and independently interesting, response is to consider ways to make the HK model more realistic. One way is to endow the agents in the model with something like a brain that is capable of learning much in the manner in which we humans learn. This is the approach to be taken here.

2.2 Multilayer perceptrons

The “brains” with which we are going to equip the agents are going to be ANNs, specifically MLPs. Or rather, our agents are going to *be* MLPs, where these MLPs form communities and attend both to worldly evidence and to their peers.

ANNs not only have a brain-inspired architecture (Goodfellow, Bengio, and Courville 2016, chap. 1); they also reflect, to some extent, how the human brain operates (Caucheteux and King 2022; Glorot, Bordes, and Bengio 2011; Goldstein et al. 2021). More important for present purposes, ANNs have been shown to be an adequate tool for simulating various higher-level cognitive processes, such as categorization, language learning, and reasoning (Battleday, Peterson, and Griffiths 2021; Buckner 2018, 2023; Douven, forthcoming; Hoffman, McClelland, and Lambon Ralph 2018; Hosseini et al. 2024).

MLPs are a specific type of ANN, belonging to the family of feed-forward ANNs, which are characterized by the unidirectional flow of data through the network. The MLP architecture dates back to the late 1950s and early 1960s, but it was not until the 1980s, with the introduction of the back-propagation algorithm (Rumelhart, Hinton, and Williams 1986), that MLPs became able to learn from complex data patterns.

An MLP consists of several densely connected layers: an input layer, one or more hidden layers, and an output layer. Each layer is made up of nodes or neurons, with all neurons in the hidden and output layers being characterized by their weights (one for each neuron in the previous layer) and biases as well as by their activation function. The weights and biases are the adjustable parameters of the network, governing the strength of connections and the threshold for neuron activation, respectively, and the activation function is typically a nonlinear function, such as a sigmoid function or rectified linear unit (ReLU) function (this was used in the studies to be reported), which calculates the neuron’s output on the basis of its inputs and the weights and bias associated with it.

MLPs are trained using a supervised learning technique (i.e., on the basis of labeled data) and learn through the aforementioned back-propagation algorithm. Specifically, the learning process in an MLP involves two phases, namely, propagation and weight updating. During propagation, unlabeled input data are passed through the network, and each neuron processes the incoming data to produce an output based on its associated weights, bias, and activation function. This output is then

passed on to the next layer until it reaches the output layer, producing the network's prediction. Next, this prediction is compared to what the output *should* have been (i.e., the label that was not provided as input). The discrepancy between prediction and label is the "error" or "loss" from the output layer, which serves as input for the back-propagation algorithm. In back-propagation, the network adjusts its weights and biases to minimize the error between its predictions and the actual outcomes. This involves calculating the gradient of the loss function with respect to each parameter (weights and biases). The network then updates its weights and biases, using gradient descent or a similar optimization algorithm, to improve its performance.

MLPs have been used for a variety of tasks (e.g., image and speech recognition, natural language processing, and time series prediction), and they have found application in a great number of areas, including financial forecasting and medical prognosis, where they aid in uncovering patterns and relationships in data that are not readily apparent to the human eye. In one of our studies, the agents (i.e., MLPs) engage in a multiclass classification task; in the other study, they are employed in the context of medical diagnostics.

3 HK updating for neural networks

The key elements of the HK model are the notion of peerhood, regulated by the ε parameter, and the operation of mixing worldly and social information, where the exact mixture depends on the value of the α parameter. We want to retain both elements in our new model, but, given that our agents are going to be MLPs, these elements need to be adapted.

In the original HK model, every agent is, at every point in time, fully characterized by its estimate of τ , making a definition of peerhood in terms of similarity of opinion the only plausible option. Accordingly, the model lets agents i and j be each other's peers at t precisely if $|x_i(t) - x_j(t)| \leq \varepsilon$. But with MLPs as agents, one agent can be similar to another agent in more than one respect. Most notably, while in the original HK model, there is no meaningful distinction between an agent's *state* at a given point in time and its *output* (i.e., its estimate of τ) at that point in time, in the new model, there is. At any point in time, an agent is in a certain state, fully characterized by its parameters (its weights and biases) at that time (architecture, including activation functions, will always be the same for all agents in a community), but it can *also* be characterized by its output (i.e., the predictions it would then make, if prompted). As a result, we can distinguish between *state-based similarity* and *output-based similarity*—and of course, agents can be similar to each other in both respects at the same time, which would make them *state- and output-based similar*.

Because MLPs can be used for various purposes, the output of an MLP can be many things: a single number, as in the HK model, or a grouping of items of interest into different classes (if the MLP is a classifier), or an assignment of probabilities to a set of competing hypotheses (e.g., if the MLP is used for a multinomial regression task), and so on. How to make the notion of output-based similarity precise will depend on the type of output we are dealing with. If, for instance, it is a single number, output-based similarity could again be defined in terms of absolute difference, as in the HK model. If the MLP is a classifier, one can consider a number of different metrics of classification similarity, such as the mutual information index, which we will use in the first study




Figure 1. Multilayer perceptrons sharing the same architecture but with different weights and biases. (Weights are annotated on the edges connecting the neurons; biases appear inside the neurons).

to be reported in this article; if the output is a probability distribution, then again, a number of options are available, such as the Kullback–Leibler (KL) divergence or the Jensen–Shannon (JS) divergence, the latter of which we will use in the second study; and so on. Two agents will then be said to be each other’s *output-based peer* precisely if they are close enough to each other in terms of the appropriate criterion.

The notion of state-based similarity requires more explanation. Given that, in our models, all agents (i.e., MLPs) will have the same architecture—the same number of layers, corresponding layers having the same number of nodes, corresponding nodes having the same activation function—we can measure their similarity by comparing their parameters, node per node. A common metric for this purpose is the cosine similarity, which requires that we vectorize the parameters first.² Gathering, in some order, the weights and biases of agent *i* in a vector params_i and proceeding analogously for the weights and biases of agent *j*, obtaining params_j , their cosine similarity is calculated as

$$\text{cossim}(i, j) = \frac{\text{params}_i \cdot \text{params}_j}{\|\text{params}_i\| \times \|\text{params}_j\|},$$

which will be a value between -1 and 1 , with 1 indicating maximum similarity and -1 maximum dissimilarity.

To illustrate, consider the MLPs shown in figure 1. We lay out sequentially the parameters of each network from top to bottom and from left to right, and we calculate the dot product of the resulting vectors:

$$0.11 \cdot 0.87 + 1.26 \cdot 0.64 + 0.43 \cdot 1.22 + \dots + 0.94 \cdot 0.81 \approx 7.35.$$

We further calculate that the norm of the first vector equals

$$\sqrt{0.11^2 + 1.26^2 + 0.43^2 + \dots + 0.94^2} \approx 2.94$$

and that the norm of the second vector equals

$$\sqrt{0.87^2 + 0.64^2 + 1.22^2 + \dots + 0.81^2} \approx 3.09.$$

Thus the cosine similarity for the aforementioned MLPs equals (approximately) $7.35 / (2.94 \times 3.09) \approx 0.81$.

² To compare networks with different architectures, other metrics than the cosine similarity are recommended (see Chen et al. 2021).




Figure 2. Multilayer perceptron with weights and biases resulting from averaging the corresponding weights and biases from the multilayer perceptrons shown in figure 1.

We will say that agents i and j are *state-based peers* precisely if $\text{cossim}(i, j) \geq 1 - \varepsilon$, for the chosen $\varepsilon \in [0, 1]$; so, for instance, if $\varepsilon = 0.2$, then the agents in the preceding illustration are each other's peer. Note that, as in the original HK model, a larger value of ε means a more liberal or inclusive notion of peerhood, which does not require agents to be as similar with respect to their parameters to qualify as peers; conversely, the smaller the value of ε is, the more similar the agents have to be, with the limiting case of $\varepsilon = 0$ meaning that the agents must be maximally similar, also analogous to the original HK model.

It merits emphasis that being able to differentiate types of peerhood—based on state, outcome, or a combination of the two—is already an enrichment compared to the original HK model. For, as social scientists have shown (e.g., Eysenbach et al. 2004; Laninga-Wijnen and Veenstra 2023), in real life, peer selection is influenced by a multitude of criteria: we may want to team up with people who share our views but also with people who look like us or have the same educational background or socioeconomic status. State-based peers could be regarded as corresponding somewhat to peers who “look like us,” output-based peers as corresponding to peers who “have views like ours.”

The averaging operation can take different forms as well, again because, with MLPs as agents, we can make a state–output distinction. Supposing we have determined an agent's peers at a given point in time (be these state-based, output-based, or state- and output-based peers), one plausible option is to average the parameters of those peers and calculate the output of the network with the resulting averages as parameters, given the input at the point in time; another, equally plausible option is to calculate the outputs of all peers at the point in time and average those outputs. In general, the results will be different. Suppose, for instance, that the two MLPs depicted in figure 1 are both given as input the vector $(2/3, 1/3)$. Then it is an easy (if somewhat tedious) exercise to calculate that the left MLP will give as output (approximately) 3.70 and the right one will give as output (approximately) 5.33, yielding an average of (approximately) 4.52. But applying the procedure of the first option to the same MLPs results in the network shown in figure 2, and this network yields (approximately) 4.46 when given $(2/3, 1/3)$ as input.³

The different definitions of peerhood and averaging can be combined in a variety of ways to obtain NABMs whose agents update in a HK-like fashion. We will make no attempt to be exhaustive here and confine ourselves to studying three models that

³ We are assuming ReLU activation functions here.

could all be plausibly regarded as extensions of the HK model, the main difference in all three cases being that the traditional HK agents have been replaced by MLPs. Roughly, the first model assumes a state-based notion of peerhood and also averages agent parameters instead of outputs. The second model assumes an output-based notion of peerhood and averages outputs. And the third model combines the first and second, which means that it proceeds by averaging parameters of state-based peers but also by averaging outputs of output-based peers. In the remainder of this section, we describe each of the models in more detail, and in the next two sections, we use computer simulations to compare their performance on standard machine learning benchmarks.

All three models require as input a community of agents, which will be MLPs but could also be different types of networks; input data, split into a training and a test set; and values for parameters regulating peerhood and the mixing of worldly and social factors in updating. There is no restriction on the exact architecture of the MLPs, except that (1) it must be the same for all MLPs in a given community, meaning that they must have the same number of layers and that corresponding layers must have the same number of nodes as well as the same activation function, and (2) the input and output layers must (of course) fit the data and task, respectively. In the first two models to be considered, the parameter ε regulates the criterion for peerhood (which, however, means different things in the two models), and the parameter α regulates the weighting of the worldly versus the social factor in updating (the weighing operation also means different things in the two models). The third model, which, as said, combines the first two, has two parameters regulating peerhood—one regulating state-based peerhood (ε_1), the other output-based peerhood (ε_2)—as well as two weight parameters, one pertaining to the weighing of states (α_1), the other pertaining to the weighing of outputs (α_2).

The first model consists of three main parts. The first part calculates a cosine similarity matrix for all agents in the community and, on that basis, selects peers for each agent (i.e., the agents that are ε -similar to it). It then calculates the averages of those parameters, in the way illustrated previously, and it stores these parameter averages. The second part, which can be thought of as the worldly part of the updating process, trains for one training round every agent (i.e., MLP) on the data it received, where it is left open at this point whether all agents receive the same data or receive different (possibly partly overlapping) subsets of the data. The third part, finally, takes a weighted average of the parameters of the agent that resulted from the training process in the second part and the parameter averages of the agent's peers that were calculated in the first part, the weighing depending on the value of α . The parameters that result from this weighted averaging are then set as the new parameters of the agent. Algorithm A.1 in appendix A.2 presents pseudo-code for the updating method defined by this model. In that presentation, the procedure outputs both the updated agents and the results from evaluating the updated agents on the relevant data (the training set, or the test set, or both, whichever is most useful for one's purposes).⁴

The building blocks of the second model are basically the same as those of the first model, but they appear in a different order. First, all agents are trained on whatever

⁴ For a still better understanding of the computational details, readers are invited to consult the Jupyter notebook in the supplementary materials, which contains the Julia code of the simulations reported in the following sections.

the relevant data are (where it is again left open whether all agents are trained on the same data or whether each agent receives its own data set); then they make predictions, whether for their training data, their test data, or both (e.g., if the task at hand is one of classification, they predict, after being trained, how each data point will be classified); in a next step, the peers of each agent are determined on the basis of how similar their predictions are (the similarity cutoff depending on ε); and finally, some α -weighted average of the agent's predictions after the worldly update and the averaged predictions of its peers is calculated and then evaluated. Algorithm A.2 presents the pseudo-code for the second model. As presented there, the procedure gives the result of the final step (i.e., of the weighted averaging) as output, together with the updated agents.

The notion of averaging, as it is used in the second model, requires a comment. Parameters are always numbers, and we know what it means to average numbers. So, in the first model, averaging always means taking the arithmetic average of whatever the relevant numbers are. But as already explained, given the many kinds of tasks MLPs can fulfill, the outputs in the case of the second model need not be numeric. As a result, the operations of averaging and weighted averaging, as carried out in the model, may differ, depending on the nature of the data or the task at hand. For instance, in the first study, averaging consists in determining the modal responses for the various data points, in a way to be detailed. Nevertheless, the intended meaning of averaging in this algorithm should be clear: the average is always some kind of best compromise of whatever different responses are under consideration.

The third model combines the previous two. Specifically, it proceeds as follows: (1) for each agent, select its state-based peers (depending on ε_1) and take the averages of their parameters; (2) train all agents on their training data; (3) for each agent, set its parameters equal to a weighted average of its parameters after the training and the averaged parameters of their state-based peers (before training; the weighting of the average is determined by α_1); (4) let all agents make predictions on the relevant data; (5) for each agent, select its output-based peers (based on ε_2), in light of the predictions obtained in the previous step; and finally, (6) for each agent, take an α_2 -weighted average of its own predictions and the averaged predictions of its output-based peers.

While in the description of the models we have not explicitly referred to an equivalent of the parameter τ in the HK model, the references to agents making predictions and being evaluated all refer *implicitly* to such an equivalent, that is, a target that the agents are aiming at and can get right to differing degrees; also, the data will, ideally, be informative of that equivalent, meaning that they will help the agent approximate the target, or even hit it. But precisely because MLPs can be used for a variety of purposes, it is impossible to characterize the target generally. If the MLPs are trained on a classification task, the aim is to classify correctly whatever data they are given as input; their predictions concern the classification of those data—they are their best guess of how the data are classified in reality—and they are evaluated in light of how closely their predictions match the correct classification, and similarly if the MLPs are trained to assign probabilities to a set of rival theories, or to predict time series, or to encrypt data. In all those cases, there is a target to which they are trying to come as close as possible and with respect to which they can

be evaluated, but the nature of the target is different each time, unlike in the HK model, where it is always a number (or a set of numbers, in some extensions).

It is again to be noted that we are not aiming at exploring all possibilities of integrating the HK model with ANNs. We do believe, however, that the three models defined in the foregoing paragraphs are all natural extensions of that model as originally conceived, the new characteristic element being that the communities of agents are constituted by neural networks. While it has been shown that, depending on an agent's environment and its goals (epistemic or otherwise) in that environment, social updating in the manner of the HK model can have notable epistemic advantages (Crosscombe and Lawry 2016; Douven 2010, 2019; Douven and Hegselmann 2021, 2022; Glass and Glass 2021), it remains to be seen whether there is any merit to HK updating for agents conceived as neural networks. To find out, the next two sections test the three models on tasks for which neural networks have been commonly used, and we compare the performance of networks in the models with that of neural networks carrying out the same tasks in a strictly individual fashion.

4 Study I: Classifying colors

The first study considers communities of agents (i.e., MLPs) that are trained to classify colors on the basis of their coordinates in color similarity space, specifically CIELUV space (see figure A.1 [left] in appendix A.3; also see, for theoretical background, Fairchild 2013). Both the training and the testing materials come from the 320 chromatic Munsell chips that served as the materials for the World Color Survey (WCS; Cook, Kay, and Regier 2005), a large catalog of color-naming systems from across the globe; the 320 chips are highlighted figure A.1 (right) and shown in a chart in the way they were presented in the WCS in figure A.2.

Because a significant number of participants in color-naming studies for both English and French used only ten of the eleven basic color terms (“green,” “blue,” etc.) in describing the colors of the WCS chips, leaving out “gray” (Berlin and Kay 1969; Claidière, Jaissati, and Chevallier 2008), we take as the target classification that the agents should try to learn—our τ , so to speak—a clustering of the WCS chips into ten categories. Also, because in the same color-naming studies, there was considerable interpersonal variability in how these chips were named, we use the *k*-means clustering algorithm to provide a kind of objective approximation of the natural color concepts.⁵ The result, which is the classification the agents should try to learn, is shown in figure A.3 (top).

The MLPs that populate the models are not much more complicated than the ones used in our earlier illustration. They also have only one hidden layer, now consisting of nine nodes and integrating the ReLU activation function for each node.⁶ Given that the task at hand is to categorize colors as belonging to one of ten classes on the basis of their CIELUV coordinates, the input layers of the MLPs have three nodes—one for

⁵ See Douven (2017, 2023) for more on this; how close the approximation is is unimportant for present purposes.

⁶ In light of recent work on ANNs, this is an exceedingly simple and shallow architecture that, by today's standards, does not even qualify as *deep* (see, e.g., Buckner 2023, 50; Buckner and Garson 2018). But everything written in this article generalizes to MLPs with any number of hidden layers and even, with some qualifications, to more recent architectures (see section 6).

each coordinate—and the output layers ten, each representing one of the basic colors minus gray.

We ran three sets of simulations, one for each of the models defined in the previous section, in which the communities always consisted of fifty MLPs with the architecture described earlier. Each simulation involved training the agents over one hundred epochs, where an epoch is a single application of the given model, with the agents being returned after epoch n serving as input for the model in epoch $n + 1$, for $n \in \{1, \dots, 99\}$.⁷ The training used the Adam optimization algorithm (with a learning rate set to 0.001) and the multiclass cross-entropy loss, which computes the loss by measuring the difference between the predicted classification probabilities (i.e., the probability that a chip should be classified as green, the probability that it should be classified as blue, etc.) and the true class labels.

Per epoch, the agents received a fresh batch of training data, each time sampled randomly and for each agent individually from the WCS chips in such a way that the number of chips from each category according to the target classification was greater than 0 but otherwise random. Thus every agent was assigned at the beginning of each epoch a set $\{(L_c^*, u_c^*, v_c^*), C_c\}_{c \in 8}$ of pairs as training data, with each pair comprising the CIELUV coordinates of some WCS chip c in sample s as well as its label C_c indicating the color it has according to the target classification.⁸ The test data on which the agents were evaluated after each epoch were always the same for all agents and consisted of the coordinates of all 320 color chips together with their labeling according to the target classification.

The evaluation used the mutual information index, which measures the similarity of different classifications (see, for why this measure is preferable to alternative measures, such as the Rand index, Pfister, Leibbrandt, and Powers 2009).⁹ To be more precise, after each epoch, we measured the accuracy of each agent by calculating the mutual information between how it classified the 320 chips in our materials and how these chips ought to be classified according to the target classification.

For the first two models, which have only one ε and one α parameter, we used a grid search strategy to approximate optimal combinations of parameters. For each combination resulting from letting ε and α range independently over the unit interval in steps of 0.025, we ran one hundred simulations as described earlier. For state-based social updating, the parameter setting $\alpha = \varepsilon = 0.9$ yielded, on average, the highest mutual information at the end of the training process. For output-based social updating, the combination of $\alpha = 0.1$ and $\varepsilon = 0.3$ did best. For the combined social updating method, which has four parameters, a grid search would have been computationally too costly, and therefore we ran a random search procedure to approximate the best setting (or a best setting; uniqueness is not guaranteed).

⁷ At start time (i.e., epoch 1), the layers of the agents were initialized using the Xavier method introduced by Glorot and Bengio (2010).

⁸ For a more detailed description of the procedure, see Douven (2023). As noted in that paper, there is no fixed sample size in this procedure, given that a *random* number of chips is sampled from each color category. The average sample size was empirically determined to be 165.02 (± 31.98).

⁹ We used the normalized version of this measure so that mutual information values were always between 0 and 1, with 0 indicating that the classifications are maximally dissimilar and 1 indicating that the classifications are maximally similar (i.e., identical). For a formal definition of this measure, and for formal definitions of all other technical notions to be used in the following discussion, see appendix A.1.

Specifically, we ran one hundred simulations for five hundred combinations of random choices (all uniformly sampled from the unit interval) for the two α and the two ε parameters, finding that the best score (i.e., the highest average mutual information) after one hundred epochs was obtained for the setting $\alpha_1 = 0.81$, $\varepsilon_1 = 0.9$, $\alpha_2 = 0.14$, and $\varepsilon_2 = 0.07$.¹⁰

To get a first impression of the accuracy that can be achieved using the different updating mechanisms with their optimal parameter setting, we trained a community of fifty agents for one thousand epochs for each of these mechanisms and compared the resulting modal classifications with the target classification. (A modal classification is the classification that gives, for each chip in our materials, the modal—that is, the most frequent—response for that chip in the given community.) For completeness, we included in the comparison the modal classification obtained from a community of fifty agents (MLPs) that do not engage in any social updating but are individually trained in the exact same way as the agents in the communities of social updaters are. It turned out that, first, the different updating methods led to modal classifications that looked almost the same and, second, that those classifications were almost identical with the target classification (see figure A.3). Indeed, a comparison with the target classification yielded the same high mutual information of 0.97 for each modal classification.

Should we conclude that the various forms of social updating are equally good but also that social updating, in whichever form, is not worth the extra effort of averaging (whether parameters or predictions, let alone both)? That would be rash, because the modal classifications tell a very incomplete story, for two reasons. First, note that modal responses can be the same even if, for one algorithm, only a small fraction of agents got the label right at the end (but wrong responses were all over the place), while for another, all, or almost all, agents got it right. Second, we will want to look at more than the end state of the training process and will also be interested in how *fast* the agents were able to learn. Perhaps all updating methods led to an excellent classification eventually, but if one already got the classification more or less right quite early on in the training process, while the other updating methods did not, then for many practical purposes, that will make the former preferable.

On these issues, figure 3 offers some helpful insights. For each of the four communities of agents under consideration (i.e., the community of nonsocial updaters and the three communities of social updaters, each using a different updating method with their optimal setting) and for each epoch, figure 3 shows the mean mutual information obtained by the agents, together with 95 percent confidence bands. We see that the combined state-based and output-based procedure swiftly surpasses the others, maintaining its lead throughout the training process.

We conducted one-way analyses of variance (ANOVAs) for the mutual information scores of the four groups after each epoch. The ω^2 -values for each ANOVA are plotted on the alternative y-axis of figure 3, a green marker indicating that the ANOVA showed group means to be significantly different, a red marker that they were *not* significantly different. An ω^2 -value greater than 0.14 is conventionally taken to indicate a large effect size, meaning in our case that, although the *modal* classifications of all communities were equally good at the end, even at the end, there were large

¹⁰ See the supplementary materials for details and additional analyses.

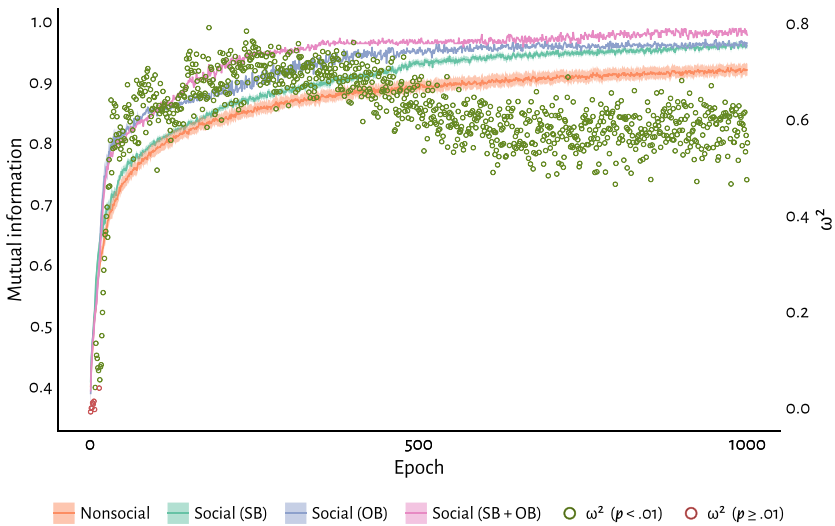


Figure 3. Per-epoch average mutual information (with 95 percent bootstrap confidence intervals) for the four communities of agents (social updating always with optimal settings; see the text). Effect sizes (ω^2) for the ANOVAs that were run for each epoch are shown on the alternative y-axis. SB, state based; OB, output based. (Color online.)

differences among the communities in terms of average mutual information scores (i.e., how well, on average, members of the communities did with respect to approximating the target classification).

Results of the per-epoch follow-up tests with pairwise comparisons, which are contained in the supplementary materials, further reveal that not only does the combined updating method top all of its rivals after virtually all epochs but at almost every epoch, a choice of the former over any of the alternatives would largely impact the achieved accuracy (where the effect size was measured using Cohen's d). The only method that at times comes close and sporadically even does better is the output-based social updating method. The pairwise comparisons also confirm what could already be guessed on the basis of figure 3, namely, that all social updating methods outperform individual updating by far.

We can also measure the total accuracy achieved by the agents over the one thousand epochs by using the area under the learning curve (AULC; see, e.g., Bouckaert 2006; Tsai, Ho, and Lin 2010), which plots the learning curve of a neural network and measures the area under that curve. Networks that learn faster and achieve greater accuracy sooner will have a larger area under the learning curve, while models that learn more slowly or achieve a lower level of accuracy will have a smaller area, assuming the same number of epochs. Thus the AULC can be interpreted as a measure of the overall performance of the network throughout the training process, with larger values for this metric indicating better average performance of a network throughout the training process.

The AULC values obtained for the agents in the community using the combined social updating method were significantly larger than those obtained for the agents in

the other communities. Specifically, a one-way ANOVA showed that type of updating had a significant and very large effect on AULC values over 1,000 epochs, $F(3, 196) = 200.58$, $p < 0.0001$, $\omega^2 = 0.75$. Pairwise t -tests showed that the AULC values for the agents using combined social updating significantly exceeded those for the agents using any other method of updating (smallest $t = 52.04$, all $ps < 0.0001$), with a mean AULC value for the combined method of $938.40 (\pm 0.56)$, for the output-based method of $921.54 (\pm 2.22)$, for the state-based method of $897.00 (\pm 0.90)$, and for the nonsocial method of $866.95 (\pm 31.44)$. A Cohen's d test showed that using the best method (combined updating) instead of its closest competitor (output-based updating) still has a large impact on overall accuracy ($d = 10.41$).

5 Study 2: Staging hypertension

The setup of the second study is broadly the same as that of the first. We look at communities of agents that aim at a target and receive evidence relevant to that target. In this study, too, the communities use different update methods, with one community consisting of nonsocial updaters and three communities consisting of social updaters, one for each of the three models from section 3.

The target is different in this study. For ease of interpretation, imagine the agents to be medical interns tasked with predicting the stages of hypertension in patients based on a variety of demographic and lifestyle data, intentionally excluding direct blood pressure readings. For the training process, we use data sourced from the National Health and Nutrition Examination Survey (NHANES), which is a yearly survey conducted by the National Center for Health Statistics.¹¹ From an initial cohort of 613 patients, we focus on the 587 adults aged twenty years and older. As key variables for analysis, we include age, gender, body mass index (BMI), diabetic status, physical activity, alcohol consumption, and smoking behavior. These variables present a mix of continuous (such as age, BMI, and physical activity), binary (gender), and ordinal (diabetic status, alcohol use, smoking) types. On the basis of these variables, the interns are to predict class probabilities for hypertension stages, ranging from normotensive (i.e., normal blood pressure), via prehypertensive and stages 1 and 2 hypertensive, to hypertensive crisis, thus encompassing five distinct categories. Hypertension stages were determined using the systolic and diastolic blood pressure readings included in the NHANES data set.

The agents are again modeled as MLPs, now comprising two hidden layers with thirty-two and sixteen nodes, respectively, employing the ReLU activation function. The input layer is designed to match the seven input variables (age, gender, etc.), while the output layer consists of five nodes corresponding to the hypertension stages. We use the softmax function in the output layer to model the output as a probability distribution, ensuring that the nodes' outputs all lie between 0 and 1 (inclusive) and that their sum equals 1.

In the training process, an agent processes data from one patient at a time and assigns, on the basis of these data, probabilities to each of the relevant hypotheses (i.e., that the patient is normotensive, that she is prehypertensive, etc.). As in the first

¹¹ We used the presently most recent batch of data available on the NHANES website, namely, the data collected from the beginning of 2017 until March 2020. The data can be downloaded from <https://www.cdc.gov/nchs/nhanes/continuousnhanes/default.aspx?Cycle=2017-2020>.

study, the training uses multiclass cross-entropy loss and the Adam optimization algorithm (with a learning rate of 0.005). This process represents the worldly part of the update, which for one community of agents is all the updating in which they engage. Three other communities of agents also participate in social updating, each using a distinct update method introduced in section 3.

As noted in section 3, state-based peers are always selected on the basis of the same criterion—namely, similarity of weights and biases—but the criterion on whose basis output-based peers are selected depends on the type of output generated. In the present case, the output consists of probability functions, and so we need a similarity measure for such functions. A prominent one is the KL divergence, but here we use the Jensen–Shannon (JS) divergence, D_{JS} , which is based on the KL divergence but which, unlike the latter, is symmetric, bounded, and normalized. Not only does that make it easier to interpret (0 indicates that the probability functions are identical, 1 that they are maximally different) but it can also be bounded by an ε parameter whose value lies in the unit interval. Thus, where one agent’s predicted probabilities at a given point in time are represented by p and another’s by q , they will be said to be each other’s peer at that time precisely if $D_{JS}(p \parallel q) < \varepsilon$, for some specified $\varepsilon \in [0, 1]$.

Relatedly, and as also previously explained, averaging of outputs will mean different things depending on the nature of the outputs. Here they are probability functions, and we use the best-known method for averaging such functions, which is linear pooling. Given probability functions $\{f_i\}_{i=1}^n$, the weighted linear average of these functions is defined to be $\sum_{i=1}^n \omega_i f_i$, with $\omega_i \geq 0$ for all i and $\sum_{i=1}^n \omega_i = 1$ (see, e.g., Dietrich and List 2016). In our model, peers are always weighted equally, meaning that we always take the straight average of their probability functions.

After each update, the agents are evaluated on the patients in the test set. Because they are making probabilistic predictions about these patients, and given that the hypotheses to which the probabilities get assigned are *ordered* (e.g., stage 1 hypertension is closer to stage 2 hypertension than the prehypertensive stage is), we evaluate the agents using the ranked probability score (RPS; see Epstein 1969). This scoring rule is particularly suited for the kind of case at hand, given that it penalizes predictions not only on the basis of how much they differ from the objective probabilities but also on the basis of the “distance” between the hypotheses in terms of their order. For example, if an agent incorrectly assigns a high probability to a stage that is adjacent to the true stage (e.g., assigning a high probability to the hypothesis that the patient has stage 1 hypertension when the patient actually has stage 2 hypertension), this is considered a less severe error than assigning a high probability to a more distant stage (e.g., assigning a high probability that the patient is normotensive, in the same scenario). After each update, we calculate the RPS for each agent and each patient, then average over all patients in the test set to obtain the overall score for the given agent after the given update. Note that lower RPS values indicate better predictive performance, with 0 being the ideal score, indicating perfect predictions.

We are interested in ascertaining whether social learning enhances the predictive accuracy of the agents and, if so, which of the social learning methods introduced in section 3 proves most effective. To optimize the parameters for the social methods, we proceed as in the previous study, performing grid searches for the state-based and output-based methods and a random search for the combined method. This yields a

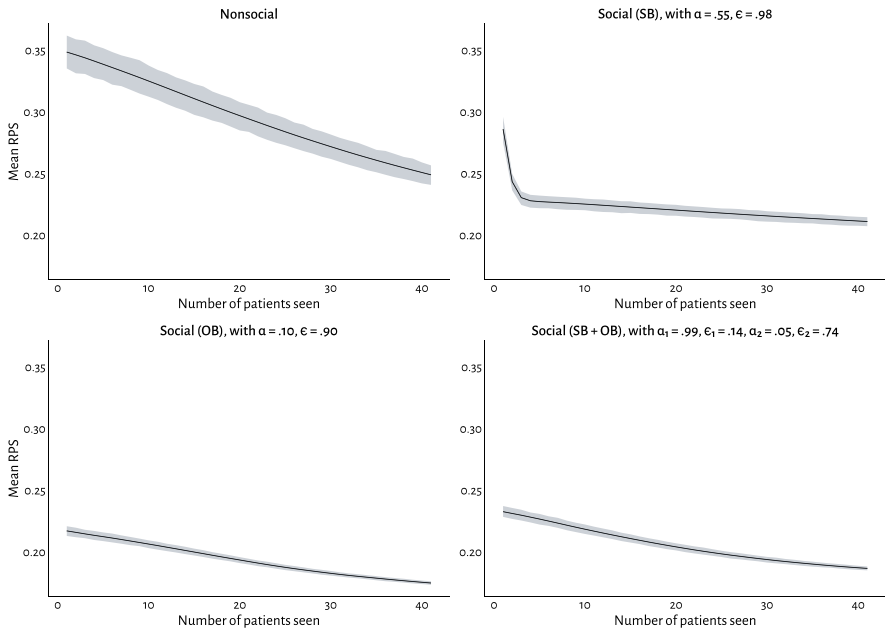


Figure 4. Per-update average (with 95 percent bootstrap confidence intervals) over fifty simulations of mean RPS achieved by agents, shown separately for the four communities of ten agents. (See the text for further explanation).

best setting of $\alpha = 0.55$ and $\epsilon = 0.98$ for the state-based method, of $\alpha = 0.1$ and $\epsilon = 0.9$ for the output-based method, and of $\alpha_1 = 0.99$, $\epsilon_1 = 0.14$, $\alpha_2 = 0.05$, and $\epsilon_2 = 0.74$ for the combined method. (See the supplementary materials for details.)

We use computer simulations to compare the social methods with optimal parameter settings both with each other and with individual updating. More specifically, we run fifty simulations, each of which starts by randomly splitting the selected NHANES data 70–30 into a training set of 410 patients and a test set of 177 patients. The 410 patients in the training set are further randomly partitioned into ten equally sized parts of forty-one patients. Each of these parts then serves as the training set of one of the interns in each of four communities of ten interns, where each community uses a different one of the four update methods in which we are interested (i.e., either individual updating or one of the three social methods).

Figure 4 shows, for the four types of communities and for each update, the average (averaged over the fifty simulations) of the average (averaged over the ten agents in the given community) RPS scored at the given update. As is already clear from the graphs, the individual updaters do, on average, worst, even by a wide margin (certainly when compared with the output-based and combined social updaters). It is equally clear that the output-based and combined social methods do better than the state-based social method. Although less clear, it seems that the output-based method does, at least for most of the updates, slightly better than the combined method.

All of this is confirmed by the ANOVAs with post hoc *t*-tests that we conducted for the simulation results per update. The outcomes are reported in the supplementary

materials, which show, among other things, that the ANOVAs were all highly significant and that they all had ω^2 -values well above 0.16. (As mentioned earlier, values for this statistic above 0.14 indicate a large effect size.) These outcomes were to be expected in light of figure 4, given the notable differences between, on one hand, the nonsocial and, on the other, all of the social updating methods. The results from the Cohen's d -tests that were also part of the follow-up tests are more informative and show that choosing a social method over nonsocial updating mostly has a large ($d > 0.8$), and always at least a medium ($d > 0.5$), impact on accuracy. And choosing the output-based method over either of the other social methods has at least, ultimately, a medium impact on accuracy. (See, again, the supplementary materials for further details.)

As we did in the first study, we end by looking at the total accuracy the agents achieved during the training process, using again the AULC. The measure of accuracy in the second study is a scoring rule that assigns penalties to agents. So, whereas in the first study, we were interested in which updating method achieved the *largest* AULC, in this study, better performance is indicated by a *smaller* area under the learning curve. A one-way ANOVA reveals a significant and substantial effect of the updating method on the accuracy of predictions; $F(3, 1996) = 189.73$, $p < 0.0001$, $\omega^2 = 0.22$. Follow-up t -tests confirm that all types of social updaters achieved significantly greater accuracy than individual updaters, which achieved a mean AULC of 11.61 (± 5.02 ; smallest $t = 11.71$, all $ps < 0.0001$, smallest $d = 0.68$). Furthermore, the state-based method users, which achieved a mean AULC of 8.74 (± 2.14), did significantly worse than both the output-based method users, with a mean AULC of 7.83 (± 1.09 ; $t = 8.63$, $p < 0.0001$, $d = 0.96$), and the combined method users, with a mean AULC of 8.75 (± 2.14 ; $t = 6.56$, $p < 0.0001$, $d = 0.39$). Finally, the output-based method users did significantly better than the combined method users, though the size of the effect is small in this case ($t = 3.28$, $p < 0.005$, $d = 0.18$).

6 Conclusion

In this article, we introduced three NABMs that extend the traditional HK model by integrating MLPs. Our models go beyond the scalar opinion representation in the HK model, enabling agents to perform complex learning tasks. Not only do the agents of the new type have enhanced learning capabilities individually but they are also capable of richer social interactions, which were seen to further improve learning.

Our computational studies, focusing on the classification of Munsell color chips and probabilistic predictions about hypertension stages, demonstrated the effectiveness of these new extensions of the HK model. Agents employing social updating consistently outperformed individual learners, underscoring the value of social learning. The results also suggest task-specific nuances in the efficacy of different updating strategies, highlighting the importance of context in social learning.

The results from our computational studies not only validate our models but also help to address the criticisms directed at agent-based modeling by, for instance, Cristelli (2014), Frey and Šešelja (2018, 2020), and Borg et al. (2019). As these critics allege, it may well be true that many ABMs are too simplistic and idealized for real-world applications. We hope to have shown, however, that this *need* not be the case

and that, by equipping ABMs with ANNs, we can model realistic forms of learning and adaptation, far beyond the limitations imposed in using traditional models like the HK model. Not only that but, using the new models, we obtained results showing the efficacy of social learning, in line with previous studies, which, however, relied on models whose validity had been called into question by the aforementioned critique.

We have limited our attention to extending one specific ABM by populating it with one specific type of ANN. It would be wrong to state that our proposal generalizes swiftly to *any* kind of ABM and *any* kind of ANN. However, many ABMs are close enough to the HK model (e.g., Friedkin and Johnsen 1990; Deffuant et al. 2000; Olsson 2013) that combining them with ANNs in the manner of this article should be straightforward. As for other network architectures, the key operations of the models proposed in this article—judging similarity on the basis of state and on the basis of output and averaging states and outputs—apply as readily to, for instance, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) as they do to MLPs. Thus an obvious avenue for future research would be to study the HK model and similar ABMs with either CNNs or RNNs as agents and see how well they do in solving tasks appropriate for the type of network used (e.g., image recognition if the agents are CNNs or predicting time series data if the agents are RNNs).

More challenging follow-up research would focus on advancing the complexity of ANNs integrated within ABMs beyond that of the ones just mentioned. Recent work has shown how LLMs can be made to communicate in that one LLM’s output serves as the prompt for one or more other LLMs, and so on, recursively (Du et al. 2023). That could be the basis for developing NABMs structurally similar to, but much more powerful than, the ones studied in this article. Naturally, comparing the internal states of LLMs is not nearly as straightforward as comparing the internal states of MLPs, and measuring the similarities between the outputs of LLMs may also be harder. But there is some work on measuring the similarity of LLMs (Chen et al. 2021), and how to compare outputs is something that will have to be decided on a case-by-case basis anyhow, as we saw already for the simple MLPs we used. Supposing these hurdles can be overcome, NABMs featuring LLMs as agents may hold the potential to enhance our understanding of complex social behaviors by enabling the study of the interplay between social learning and advanced forms of reasoning (on inductive and abductive reasoning in LLMs, see, e.g., Liu, Neubig, and Andreas 2024) within agent-based simulations.

Acknowledgments. I am greatly indebted to Rainer Hegselmann, Christopher von Bülow, and two anonymous referees for valuable comments on previous versions of this article.

References

- Adam, D. 2020. “The Simulations Driving the World’s Response to COVID-19.” *Nature* 580:316–18. <https://doi.org/10.1038/d41586-020-01003-6>.
- Battleday, R. M., J. C. Peterson, and T. L. Griffiths. 2021. “From Convolutional Neural Networks to Models of Higher-Level Cognition (and Back Again).” *Annals of the New York Academy of Sciences* 1505 (1):55–78. <https://doi.org/10.1111/nyas.14666>.
- Berlin, B., and P. Kay. 1969. *Basic Color Terms*. Stanford, CA: CSLI.
- Bezanson, J., A. Edelman, S. Karpinski, and V. B. Shah. 2017. “Julia: A Fresh Approach to Numerical Computing.” *SIAM Review* 59 (1):65–98. <https://doi.org/10.1137/141000671>.

- Borg, A., D. Frey, D. Šešelja, and C. Straser. 2019. "Theory-Choice, Transient Diversity and the Efficiency of Scientific Inquiry." *European Journal for Philosophy of Science* 9:26. <https://doi.org/10.1007/s13194-019-0249-5>.
- Bouckaert, R. R. 2006. "Efficient AUC Learning Curve Calculation." In *AI 2006: Advances in Artificial Intelligence*, edited by A. Sattar and B.-H. Kang, 181–91. Berlin: Springer. https://doi.org/10.1007/11941439_20.
- Buckner, C. 2018. "Empiricism without Magic: Transformational Abstraction in Deep Convolutional Neural Networks." *Synthese* 195:5339–72. <https://doi.org/10.1007/s11229-017-1622-1>.
- Buckner, C. 2023. *From Deep Learning to Rational Machines*. Oxford: Oxford University Press.
- Buckner, C., and J. Garson. 2018. "Connectionism and Post-connectionist Models." In *The Routledge Handbook of the Computational Mind*, edited by M. Sprevak and M. Colombo, 76–91. London: Routledge.
- Caucheteux, C., and J. R. King. 2022. "Brains and Algorithms Partially Converge in Natural Language Processing." *Communications Biology* 5:134. <https://doi.org/10.1038/s42003-022-03036-1>.
- Chen, Z., Y. Lu, W. Yang, Q. Xuan, and X. Yang. 2021. "Graph-Based Similarity of Neural Network Representations." ArXiv. <https://doi.org/10.48550/arXiv.2111.11165>.
- Claidière, N., Y. Jraissati, and C. Chevallier. 2008. "A Colour Sorting Task Reveals the Limits of the Universalist/Relativist Dichotomy: Colour Categories Can Be Both Language Specific and Perceptual." *Journal of Cognition and Culture* 8 (3–4):211–33. <https://doi.org/10.1163/156853708X358260>.
- Cook, R. S., P. Kay, and T. Regier. 2005. *The World Color Survey Database: History and Use*. Amsterdam: Elsevier.
- Cristelli, M. 2014. *Complexity in Financial Markets*. Cham, Switzerland: Springer.
- Crosscombe, M., and J. Lawry. 2016. "A Model of Multi-agent Consensus for Vague and Uncertain Beliefs." *Adaptive Behavior* 24 (4):249–60. <https://doi.org/10.1177/1059712316656890>.
- Deffuant, G., D. Neau, F. Amblard, and G. Weisbuch. 2000. "Mixing Beliefs among Interacting Agents." *Advances in Complex Systems* 3 (1):87–98. <https://doi.org/10.1142/S0219525900000078>.
- Dietrich, F., and C. List. 2016. "Probabilistic Opinion Pooling." In *The Oxford Handbook of Probability and Philosophy*, edited by A. Hájek and C. Hitchcock, 519–41. Oxford: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199607617.013.22>.
- Dittmer, J. C. 2001. "Consensus Formation under Bounded Confidence." *Nonlinear Analysis* 47 (7):4615–21. [https://doi.org/10.1016/S0362-546X\(00\)00571-2](https://doi.org/10.1016/S0362-546X(00)00571-2).
- Douven, I. 2010. "Simulating Peer Disagreements." *Studies in History and Philosophy of Science, Part A* 41 (2):148–57. <https://doi.org/10.1016/j.shpsa.2010.03.002>.
- Douven, I. 2017. "Clustering Colors." *Cognitive Systems Research* 45:70–81. <https://doi.org/10.1016/j.cogsys.2017.01.001>.
- Douven, I. 2019. "Optimizing Group Learning: An Evolutionary Computing Approach." *Artificial Intelligence* 275:235–51. <https://doi.org/10.1016/j.artint.2019.06.005>.
- Douven, I. 2023. "The Role of Naturalness in Concept Learning: A Computational Study." *Minds and Machines* 33:695–714. <https://doi.org/10.1007/s11023-023-09652-y>.
- Douven, I. 2024. "Pandemics and Flexible Lockdowns: In Praise of Agent-Based Modeling." *European Journal for Philosophy of Science* 13:35.
- Douven, I. Forthcoming. "The Learnability of Natural Concepts." *Mind and Language*.
- Douven, I., and R. Hegselmann. 2021. "Mis- and Disinformation in a Bounded Confidence Model." *Artificial Intelligence* 291:103415. <https://doi.org/10.1016/j.artint.2020.103415>.
- Douven, I., and R. Hegselmann. 2022. "Network Effects in a Bounded Confidence Model." *Studies in History and Philosophy of Science, Part A* 94:56–71. <https://doi.org/10.1016/j.shpsa.2022.01.001>.
- Du, Y., S. Li, A. Torralba, J. B. Tenenbaum, and I. Mordatch. 2023. "Improving Factuality and Reasoning in Language Models through Multiagent Debate." ArXiv. <https://doi.org/10.48550/arXiv.2305.14325>.
- Epstein, E. S. 1969. "A Scoring System for Probability Forecasts of Ranked Categories." *Journal of Applied Meteorology* 8 (6):985–87. [https://doi.org/10.1175/1520-0450\(1969\)008<0985:ASSFPF>2.0.CO;2](https://doi.org/10.1175/1520-0450(1969)008<0985:ASSFPF>2.0.CO;2).
- Eysenbach, G., J. Powell, M. F. Englesakis, C. Rizo, and A. Stern. 2004. "Health Related Virtual Communities and Electronic Support Groups: Systematic Review of the Effects of Online Peer to Peer Interactions." *British Medical Journal* 328:1166. <https://doi.org/10.1136/bmj.328.7449.1166>.
- Fairchild, M. D. 2013. *Color Appearance Models*. Hoboken, NJ: John Wiley.

- Frey, D., and D. Šešelja. 2018. "What Is the Epistemic Function of Highly Idealized Agent-Based Models of Scientific Inquiry?" *Philosophy of the Social Sciences* 48 (4):407–33. <https://doi.org/10.1177/0048393118778564>.
- Frey, D., and D. Šešelja. 2020. "Robustness and Idealizations in Agent-Based Models of Scientific Interaction." *British Journal for the Philosophy of Science* 71 (4):1411–37. <https://doi.org/10.1093/bjps/axy073>.
- Friedkin, N. E., and E. C. Johnsen. 1990. "Social Influence and Opinions." *Journal of Mathematical Sociology* 15 (3–4):193–206.
- Glass, C., and D. H. Glass. 2021. "Opinion Dynamics of Social Learning with a Conflicting Source." *Physica A: Statistical Mechanics and Its Applications* 563:125480. <https://doi.org/10.1016/j.physa.2020.125480>.
- Glorot, X., and Y. Bengio. 2010. "Understanding the Difficulty of Training Deep Feedforward Neural Networks." In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 249–56. N.p.: ML Research Press. <http://proceedings.mlr.press/v9/glorot10a.html>.
- Glorot, X., A. Bordes, and Y. Bengio. 2011. "Deep Sparse Rectifier Neural Networks." In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 315–23. N.p.: ML Research Press. <http://proceedings.mlr.press/v15/glorot11a.html>.
- Goldstein, A., Z. Zada, E. Buchnik, M. Schain, A. Price, B. Aubrey, S. A. Nastase et al. 2021. "Thinking Ahead: Spontaneous Prediction in Context as a Keystone of Language in Humans and Machines." *BioRxiv*. <https://doi.org/10.1101/2020.12.02.403477>.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.
- Hegselmann, R. 2023. "Bounded Confidence Revisited: What We Overlooked, Underestimated, and Got Wrong." *Journal of Artificial Societies and Social Simulation* 26 (4):11–15. <https://doi.org/10.18564/jasss.5257>.
- Hegselmann, R., S. König, S. Kurz, C. Niemann, and J. Rambau. 2015. "Optimal Opinion Control: The Campaign Problem." *Journal of Artificial Societies and Social Simulation* 18 (3):18. <http://jasss.soc.surrey.ac.uk/18/3/18.html>.
- Hegselmann, R., and U. Krause. 2002. "Opinion Dynamics and Bounded Confidence: Models, Analysis, and Simulations." *Journal of Artificial Societies and Social Simulation* 5 (3):2. <http://jasss.soc.surrey.ac.uk/5/3/2.html>.
- Hegselmann, R., and U. Krause. 2006. "Truth and Cognitive Division of Labor: First Steps towards a Computer Aided Social Epistemology." *Journal of Artificial Societies and Social Simulation* 9 (3):10. <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- Hegselmann, R., and U. Krause. 2015. "Opinion Dynamics under the Influence of Radical Groups, Charismatic Leaders, and Other Constant Signals: A Simple Unifying Model." *Networks and Heterogeneous Media* 10 (3):477–509. <https://doi.org/10.3934/nhm.2015.10.477>.
- Hegselmann, R., and U. Krause. 2019. "Consensus and Fragmentation of Opinions with a Focus on Bounded Confidence." *American Mathematical Monthly* 126 (8):700–716. <https://doi.org/10.1080/00029890.2019.1608551>.
- Hoffman, P., J. L. McClelland, and M. A. Lambon Ralph. 2018. "Concepts, Control, and Context: A Connectionist Account of Normal and Disordered Semantic Cognition." *Psychological Review* 125 (3):293–328. <https://doi.org/10.1037/rev0000094>.
- Hosseini, E. A., M. Schrimpf, Y. Zhang, S. Bowman, N. Zaslavsky, and E. Fedorenko. 2024. "Artificial Neural Network Language Models Predict Human Brain Responses to Language Even after a Developmentally Realistic Amount of Training." *Neurobiology of Language* 5 (1):43–63. <https://doi.org/10.1101/2022.10.04.510681>.
- Huang, A. C. W. Forthcoming. "Track Records: A Cautionary Tale." *British Journal for the Philosophy of Science*. <https://doi.org/10.1086/728459>.
- Jacobmeier, D. 2004. "Multidimensional Consensus Model on a Barabási–Albert Network." *International Journal of Modern Physics, Part C* 16 (4):633–46. <https://doi.org/10.1142/S0129183104006750>.
- Kummerfeld, E., and K. J. S. Zollman. 2016. "Conservatism and the Scientific State of Nature." *British Journal for the Philosophy of Science* 67 (4):1057–76. <https://doi.org/10.1093/bjps/axv039>.
- Laninga-Wijnen, L., and R. Veenstra. 2023. "Peer Similarity in Adolescent Social Networks: Types of Selection and Influence, and Factors Contributing to Openness to Peer Influence." In *Encyclopedia of Child and Adolescent Health*, edited by B. Halpern-Felsher, 196–206. Oxford: Academic Press.

- Liu, E., G. Neubig, and J. Andreas. 2024. "An Incomplete Loop: Deductive, Inductive, and Abductive Learning in Large Language Models." ArXiv. <https://doi.org/10.48550/arXiv.2404.03028>.
- Lorenz, J. 2008. "Fostering Consensus in Multidimensional Continuous Opinion Dynamics under Bounded Confidence." In *Managing Complexity: Insights, Concepts, Applications*, edited by D. Helbing, 321–34. Berlin: Springer. https://doi.org/10.1007/978-3-540-68421-5_17.
- Lorig, F., E. Johansson, and P. Davidsson. 2021. "Agent-Based Social Simulation of the COVID-19 Pandemic: A Systematic Review." *Journal of Artificial Societies and Social Simulation* 24 (3):5. <https://doi.org/10.18564/jasss.4606>.
- O'Connor, C., and J. O. Weatherall. 2019. *The Misinformation Age: How False Beliefs Spread*. New Haven, CT: Yale University Press.
- Olsson, E. J. 2013. "A Bayesian Simulation Model of Group Deliberation and Polarization." In *Bayesian Argumentation*, edited by F. Zenker, 113–33. Dordrecht, Netherlands: Springer.
- Olsson, E. J., and A. Vallinder. 2013. "Norms of Assertion and Communication in Social Networks." *Synthese* 190:2557–71. <https://doi.org/10.1007/s11229-013-0262-2>.
- Pfitzer, D., R. Leibbrandt, and D. Powers. 2009. "Characterization and Evaluation of Similarity Measures of Pairs of Clusterings." *Knowledge and Information Systems* 19:361–94. <https://doi.org/10.1007/s10115-008-0145-2>.
- Pluchino, A., V. Latora, and A. Rapisarda. 2006. "Compromise and Synchronization in Opinion Dynamics." *European Physical Journal B* 50:169–76. <https://doi.org/10.1140/epjb/e2006-00190-1>.
- Rosenstock, S., C. O'Connor, and J. Bruner. 2017. "In Epistemic Networks, Is Less Really More?" *Philosophy of Science* 84 (2):234–52. <https://doi.org/10.1086/690717>.
- Rumelhart, D. E., G. E. Hinton, and R. J. Williams. 1986. "Learning Representations by Back-Propagating Errors." *Nature* 323:533–36. <https://doi.org/10.1038/323533a0>.
- Schelling, T. C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1 (2):143–86.
- Šešelja, D. 2019. "Some Lessons from Simulations of Scientific Disagreements." *Synthese* 198:6143–58. <https://doi.org/10.1007/s11229-019-02342-6>.
- Thicke, M. 2020. "Evaluating Formal Models of Science." *Journal for General Philosophy of Science* 51:315–35. <https://doi.org/10.1007/s10838-020-09496-6>.
- Tsai, M.-H., C.-H. Ho, and C.-J. Lin. 2010. "Active Learning Strategies Using SVMs." *The 2010 International Joint Conference on Neural Networks*, 1–8. New York: IEEE. <https://doi.org/10.1109/IJCNN.2010.5596668>.
- Zollman, K. J. S. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5):574–87. <https://doi.org/10.1086/525605>.
- Zollman, K. J. S. 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72:17–35. <https://doi.org/10.1007/s10670-009-9194-6>.