An Opinion Dynamics of Science? Agent-Based Modeling of Knowledge Spread

Bernardo S. Buarque[1], Malte Vogl[1], Roberto Lalli[1,2], and Dirk Wintergrün[1,3]

[1] Max Planck Institute for the History of Science, 14195 Berlin, Germany
bsbuarque@mpiwg-berlin.mpg.de
[2] DIMEAS, Politecnico di Torino, 10129 Torino, Italy
[3] Klassik Stiftung Weimar, 99423 Weimar, Germany

Abstract. We present a socio-epistemic model of science inspired by the existing literature on opinion dynamics. We place the agents into social networks and put them into an epistemic space - a three-dimensional lattice where each site represents a unique topic or concept. We arrange this space according to the similarity between issues and allow the agents to move across it. They learn from each other, explore their local lattice, and collect new thoughts and ideas about their mental representations of the world. Ultimately, we keep track of every movement by the agent across the epistemic space seeking to understand the popularity of different knowledge clusters or scientific fields. Therefore, we propose an analytical model that examines the connection between agents' accumulated knowledge, social learning, and the span of attitudes toward mental models in an artificial society. While we rely on the example from the General Theory of Relativity Renaissance, our goal is to observe what determines the creation and diffusion of mental models. We offer quantitative and inductive research, which collects data from an artificial environment to elaborate generalized theories about the evolution of science.

Keywords: Opinion Dynamics · Socio-epistemic Networks · Metascience

1 Introduction

We present a dynamic model illustrating the structural changes in the socio-epistemic networks of physics. In particular, we focus on the collaborative, citation, and epistemic transformations witnessed during the so-called Renaissance of the General Theory of Relativity. It was a period of fast shifts in the field when the theory first developed by Einstein went from marginal to a pillar of modern physics [1]. Yet, we know little about the potential causes behind this renaissance - nor why it took perhaps so long to happen. Hence, we hope our data and methods can help uncover the causes, markers, and consequences of developing scientific theories and fields.

While we focus on the General Theory of Relativity, we wish more broadly to observe what influences the slow diffusion process and consolidation of new scientific paradigms. In other words, we seek to reconstruct and examine a "scientific tipping point" as emerging from the interactions of agents under a "small number of simple assumptions" [2]. Indeed, there is ample evidence of similar "contagious" processes whereby "solitary" discoveries slowly transform into established scientific communities with shared practices and commitments [3].

Succinctly, we are working towards an agent-based model of science inspired by the existing literature on opinion dynamics. Like previous models, in this ABM, we arrange the agents into a social network and give them mental representations of the world. We allow them to interact, learn from each other, and collect new thoughts and ideas for their mental models. Therefore, the agents continually update their beliefs about the world. And we keep track of their movements seeking to understand the popularity of different mental models. We allow this artificial society to progress as we try to observe the clustering of agents around a theme or an "endogenously emerging coordination" of scientists around "shared ideas and commitments" [4].

2 Data & Methods

2.1 The Epistemic Layer

We represent each knowledge unit as a triple, where elements can take the value of any integer between 0 and 256. And we arrange these units in a three-dimensional cube according to their cognitive similarity [5]. Thus, we can imagine the epistemic layer as representing an RGB function.

Besides the three-dimensional vector that gives each site a unique identity, we track the number of agents currently located in their vicinity and the number of agents who added the site to their mental models.

We use these variables to calculate a unique fitness value for each site - i.e., their attractability to surrounding agents. We assume that if there are too few agents in the position, the probability of recognition is low. And when there are too many agents, it is harder to stand out [4]. Therefore we model fitness as a logarithmic function of popularity:

$$\Pi_{st} = \ln(n_{st}) \tag{1}$$

where n_{st} represents the site's population. We measure it as the count of agents with the site's triple in their knowledge stocks and the number of agents currently located within a given radius.

2.2 The Agents

Each agent represents a scholar. They are heterogeneous and described by static and dynamic variables. We attribute them to social ties and a mental model - i.e., a list containing all the topics they acquired. We also give the agents an activation probability and a time-varying reputation or acknowledgment accrued from all their previous publications.

Mental Models (κ_{it}) We represent mental models as hypergraphs where each subgroup is a unique combination of knowledge units common to the model. Each agent has its mental representation of the world, which developed from its historical paths across the epistemic layer - all the sites they've visited. Thus, we can interpret the mental model as a growing list containing every triple the agent accepted or learned. For example, we could represent the mental model of one agent as $[(10, 20, 10), (10, 25, 10), ..., (15, 25, 10)]$.

The Simulating Knowledge Dynamics in Innovation Networks platform serves as our inspiration [5]. But we extend the original metaphor and simulate agents building on their models by "piecing together" different bits of knowledge - like a jigsaw puzzle [6], ingredients in a recipe [7], words making up a vocabulary [8] or a network [9]. Furthermore, we represent the evolution of their models similar to the seashore walk analogy in [10].

Activation Probability (α_i) Agents have distinct propensities to engage with their surroundings - much like some authors are more prolific than others, and people differ regarding how often they communicate or interact with one another. Along these lines, the agents' activation probability informs the likelihood we will initialize the agent at each time step. The activation rate follows a power-law distribution, and at the start of the simulation, we attribute to each agent a unique value between $[0, 1]$.

Acknowledgement (π_{it}) It describes the agents' reputation or prestige. Every time the agent publishes a paper, they receive credit and acknowledgement. The value they receive for each publication comes from a normal distribution centred around the topic's fitness - i.e., the number of agents accepting the idea. And we model agents' acknowledgement as the sum of all credits or values they collected from all their publications.

We use acknowledgement as the agents' profit function in the model. Although we don't describe the agents as profit-maximizing, including the variable can lead to several possibilities for the ABM. First, we can use it in the repulsion-attraction functions representing the influence agents exert over others - much like the role of confidence in the basic opinion dynamics model [11]. Thus, differences in reputation could lead to more heightened attraction forces between a postdoc and a high-ranking professor.

Besides, we can use the variable to grow the social network. We could use a variation to preferential attachment where incoming and living agents are more likely to form connections to those with higher reputations.

2.3 The Social Layer

We connect the agents using an unweighted, undirected, and temporal network. The ties between scholars portray co-authors or those working at the same place. We assume the temporal graph follows a small-world or power-law distribution. But its underlying structure depends on the sub-model specification so that we can study how different social arrangements influence aggregate behaviour.

Population Growth (δ) To simulate the exponential growth in PhD graduates and postdoctoral researchers observed in the late 1950s [12], we include between ticks a populational growth term:

$$N_t = (1+\delta)N_{t-1} \tag{2}$$

In other words, during every tick of the model, we add agents to the system. These represent the PhD graduates, and we can think of them as "offsprings" of their supervisors. Following a preferential attachment, newborns link to potential supervisors according to the latter's reputation or acknowledgement. And we don't assume newborns enter the system as a blank slate. Instead, they partially inherit their supervisor's mental models and connections.

We model a birth-death process akin to evolutionary computational models, where acknowledgement determines the likelihood of "reproducing." Newborns randomly select a parent (or supervisor) weighted by their fitness, then they form links to their supervisors and inherit their connections and models.

In line with [15], we assume that newborns always connect to their parents; they inherit each of their parents' links with a given probability and form new connections with a different and smaller chance. Furthermore, we assume that each new tie modifies the newborn's inherited model to some extent - i.e., the new social reference pulls the newborn's model from its original form [2].

Sociability Growth Besides having more people, we observe more connections between them. The network's average degree grows in time. To account for it, we allow the nodes already present in the network to form new links following preferential attachment. Thus, we assume the probability of a new link between two existing nodes is proportional to the product of their connectivities. We calculate the likelihood of making a new connection as:

$$a_{ij}^N = \beta * \frac{k_i k_j}{\sum k_s k_m} N_t \tag{3}$$

where k_i is the degree centrality of the agent i. And the parameter β refers to the number of newly created internal links per node in unit time.

The basic design for sociability growth comes from [16] work on the evolution of scientific collaborations. They propose a simple generative model that reproduces the following empirical results: average path length decreases in time, clustering coefficient also decays, average degree increases, and the relative size of the largest cluster grow. Therefore, the benchmark reproduces all aspects of sociability growth we need, is prevalent in the relevant literature, and uses a single parameter. But we also try alternative variations. For example, instead of relying solely on connectivity, the preferential attachment could account for the agent's activation probability and acknowledgment - i.e., younger and proactive agents are more likely to start connections to older and famous scientists.

2.4 Opinion Dynamics

In each round, agents try to learn concepts from their peers, and we model this process as them moving towards each other in the epistemic layer. Namely, the agents randomly select someone from their social network and converge to their current position in the three-dimensional cube.

The "speed" at which the agent moves is a logistic function of personal attributes, social status, and epistemic distance:

$$\Delta K_i^S = \frac{1}{1 + e^{-\phi(\kappa_{it} - \kappa_{jt})}} \tag{4}$$

where the weighting parameter ϕ is a function of the ratio between the agent's and their focal point's acknowledgement. κ is their stock of ideas.

The current description follows the most basic setup for an opinion dynamics model [11]. Like most previous studies, first, the agents choose a focal individual and then partially assimilate their opinion. Many factors - e.g., confidence - can affect how much influence one agent exerts over others, and we could equally account for those. In the logistic function, we can weigh the cognitive distance between agents by their "status difference" or their "open-mindedness."

Environmental Learning Agents also move toward nearby topics according to their popularity. We model their decision as an urn process whereby the probability of selecting a site within a given radius is proportional to their fitness. After choosing a focal site, the agent moves towards it following a similar function to the one described in the social learning:

$$\Delta K_i^E = \frac{1}{1 + e^{-\lambda(\kappa_{it} - \kappa^*)}} \tag{5}$$

where the weighting parameter λ is proportional to the site's fitness and the agent's open-mindedness or activation rate. κ^* is the site's triple.

To keep it in line with social learning, we model environmental search by, first, asking the agents to find a random "focal point" to which they will move. So, in each round, they move towards one random site in their vicinity - where the attraction force and the likelihood they choose this point are both a function of the site's popularity or fitness.

Hence, the current version differs from typical models using the "search in rugged landscapes" metaphor [17]. Unlike the former, the agents do not compare the payoffs between neighbour sites before moving from a lower to a higher fitness point. The model also differs from works like [18], where they have a global attraction force based on the suitability of different sites. These serve as inspiration, but from there, we take from "gravitational models" in social sciences [19]. In particular, the idea is to reproduce past research on human mobility, which shows that people "choose a new location to visit depending on both its distance from the current position, as well as its relevance" [20].

Conferences Beyond social and environmental learning, a central premise in this model is the role of conferences. To make these, we randomly sample a number of agents - where the the group size is also random. And we find the central point between them in the epistemic layer - the centroid of their positions in the three-dimensional cube. Next, we define this centroid as the agents' focal point and request them to move towards it. The "speed of convergence" follows the same rules as in the environmental learning process. Thus, the only difference is how we find the focal point.

We can interpret the conferences as a process of "group learning." In contrast to social learning, therefore, when part of a conference, agents do not engage in pairwise interactions [21]. Instead, all the attendees must move together towards a central theme or topic. So, conferences work closer to a voter or majority rules model [22] than the typical opinion dynamics paradigm in ABM.

Mining for acknowledgment After moving to a new location, the agent strives to publish using this knowledge unit. And they collect credit from their work. We model it as a stochastic process following a normal distribution with the mean equal to the site's fitness and standard deviation proportional to the average distance between the agents' mental model and the lattice.

$$\pi_{it} \sim \mathcal{N}(\Pi,\ \langle d(\kappa, \kappa^*) \rangle) \tag{6}$$

where Π is the site's fitness and $\langle d(\kappa, \kappa^*) \rangle$ is the average distance between each unit in the agent's mental model and the site's triple.

2.5 Data & Inputs

Our goal is to study the complex socio-epistemic transformations occurring in physics post-1950s. Thus we can use data from the period before to calibrate the ABM. Borrowing from the machine learning terminology, we use the period between 1920 and 1950 akin to a training set. So, like [1], we collect data covering publications in theoretical physics from the Web of Science and NASA ADS. And we exploit it to construct our agents, attribute their initial conditions, place them into social networks, and so forth. We use it to set up our artificial society.

We can follow the career paths of scientists to measure, for example, how often they publish or the average distance between their keywords. Then we normalize their records to get a distribution of value, which we use to initialize the agents. Namely, we use this distribution to give them an activation probability proportional to how often physicists published in 1920-1950, assign acknowledgment values proportional to the number of citations, give them a list of topics, place them in the epistemic layer, etc.

In addition, we use the data gathered by [1] to build and grow social networks. Beyond looking at co-authors and co-workers, their data provide extensive information covering all ties between those working on the field. They provide information on PhD graduates and their supervisors, the conferences happening at the time and who attended them, and much more. Thus, we benefit from this wealth of information to set up our social layer.

The General Relativity Renaissance is our case study. Therefore we can employ the data from [1] to calibrate and validate our model. However, as usual, we must acknowledge here the "danger that the simulation would be trivial and reproduce the particular behavior without providing an insight into the generality of the mechanism involved" [23]. We do not want to tailor the model to the specific case of GR. Instead, we might need to have some variability regarding the initial distribution of agents and mental models - as these could contribute to our understanding of the impact of social structures on the epistemic layer.

2.6 Outputs

The keep the model as simple as possible, we start the epistemic layer as a blank slate. We construct the three-dimensional cube with the desired size and proportions. But we don't attribute values, topics, or fitness to the sites. Instead, we use the model inputs to start and run different populations moving across this blank cube. So, we treat it as a canvas, where we can see how particular social arrangements and starting positions might influence the final picture. Along these lines, the variables that make the site's popularity and fitness are the chief outputs of the model. Similar to a double-ledger, we will have a record of every agent's history - their mental models. And also an environmental chronology of how many people accepted the idea at a given time and the concentration or dispersion of agents/publications in the cube. At last, we can use this data to make colorful graphs showing the impact of the social layer on the epistemic one. Since we are using the RGB function, our agents are technically learning and discovering color palettes - so we can use it to "paint" immersive results.

2.7 Extending the Model

Building on the seashore metaphor, we will only allow the agents to move around the three-dimensional cube. At each step, they stop at a different site and collect a triple - a topic associated with their location. In other words, currently, the mental models only represent a list of sites - a collection of ideas the agent acquired at some point. But we could extend the method to allow the agents to create or develop new ideas based on those they already hold. Instead of just collecting or learning new "colours" from the environment, the agents can interact with the units they know. They will "break, blend, and bind" the colours in their stock to discover new elements [24]. The ABM would then progress similarly to the Adder model [25] - where agents learn new numbers and formulas by adding/subtracting the values in their repertoire. In sum, not only should the agents collect new tools but also constantly upgrade, update, forget, and edit their mental models - though this is beyond the focus for now.

Another matter that requires future work is how to efficiently combine the different influences on the agent - i.e., how to update the agents' position at each step. Currently, we update each agent according to each function in a sequence

- first, we move them toward their colleagues, then to a random topic, and so forth. But perhaps there are better ways to describe this process. For example, we could use a probabilistic function where each type of learning occurs with a different frequency - perhaps we can make environmental learning more common but weaker. Or we could have some agents more likely to engage in social learning than others [26].

The basic assumption behind the model is that agents explore different topics surrounding their interests but might also learn about far-removed issues from their social ties. Because we are interested in the social influence on the epistemic layer, it's essential to show, for example, that social learning and increased sociability can lead to different paths in the epistemic layer. Still, it's difficult to separate or delineate between the forces. Under the current setup, for instance, we might think of environmental learning as a form of quorum sensing, where agents decide the most relevant topics by looking at the number of agents around it. Hence, we might also accept it as a social influence between individuals - albeit an indirect effect. With that in mind, a simple yet elegant solution is to combine all the elements into "one big urn model." At each step, the agents try to learn a new, focal topic. Their ability to efficiently acquire the "color" depends on their distance and some personal attributes. But, most importantly, the likelihood they will choose the topic must rely on several factors - e.g., its popularity, how many of my friends know the issue, how often I visited the site in the past, etc. And this must be a self-reinforcing urn - we constantly update the probabilities of picking a site.

At last combining both issues discussed, we are currently working to introduce a new schedule where, first, the agents try to acquire a new color. When successful, they add the new triple to their repertoire but then engage in creative behaviour - i.e., they mix the new set with colours they acquired in the past - to produce even more new pallets. Along these lines, the goal/role of communication is to inform others about your newly found/acquired colors.

3 Discussion

Our model borrowed extensively from previous findings in social and cultural evolution. However, most other models fail to produce a similar phase transition observed during the general relativity renaissance. In a sense, we lack a robust description of the ever-changing nature of science, whereby new clusters endogenously emerge but also break - thus leading to yet new paradigms and conventions. The models we draw inspiration from do not acknowledge the inertia behind the forces that explain how "minority groups can initiate social change dynamics in the emergence of new social conventions" [27].

Succinctly, we seek an explanation for the renaissance that breaks the original order and leads the system to a new path. We need a source of external energy that breaks the inertia. And we are working with a few hypotheses and ideas. First, we consider the impact of population growth. The general relativity renaissance took place between the 1950s and 1970s. During this time, we observe a nearly threefold increase in PhDs and postdocs entering the market. So we ought to account for said growth.

Population growth is a feature of most ABMs in cultural accumulation - e.g., those dealing with the diffusion of tools in early societies [13]. Yet, they are seldom part of opinion dynamics models. Perhaps because including population growth prevents the system from stabilizing and makes analytical solutions almost impossible [14]. Altogether, we observe population growth in our data, and still, we have little knowledge about how it will influence the dynamics of our system. So we must depart from classic opinion models and include the exponential population growth parameter as a sub-model definition of the ABM.

Besides having more people, we also observe more or perhaps different connections between them. The technological progress of the time made it particularly easy to meet and maintain relations with authors far apart. And as expected, we see that the average number of co-authors and citations to others grows during this time. But we also notice new types of social connectivity or, at least, new opportunities and venues for scholars to interact. Indeed, a common feature of many developing disciplines is the development of new institutes, dedicated volumes, and special issues [28]. The specialized venues organize the social network into localized clusters with wide bridges. So, more than simply becoming more connected, the network also evolves to allow for partial fragmentation or modularity [29] [30].

Following the last argument, a central premise in this model is the role of conferences. And in the ABM, we interpret the conferences as a "dual medium" learning model. In other words, we intend to capture the duality of people learning and engaging in different environments or social layers. Like scientists who engage daily in compartmentalized laboratories and travel to attend conferences, we allow the agents in the ABM to interact with others using two distinct "rules." During "normal" circumstances, they behave like classic opinion dynamic models. They assimilate the opinions of those in their social network, one at a time. But, at different moments, and with a certain probability, the agents might attend a conference. There, they will engage in "group learning."

Along these lines, adding the conferences to the model may introduce a higher level of complexity and further diversity into the system. Borrowing again from [14] thermal analogy, it can serve as a high-temperature micro-environment that injects fluctuations away from the system's inertia and towards new social conventions - i.e., new scientific paradigms.

Acknowledgements The development is part of the research project ModelSEN Socio-epistemic Networks: Modelling Historical Knowledge Processes, part of Department I of the Max Planck Institute for the History of Science and funded by the Federal Ministry of Education and Research, Germany (Grant No. 01 UG2131).

References

1. Lalli, R., Howey, R., Wintergrün, D.: The dynamics of collaboration networks and the history of general relativity. Scientometrics **122**, 1129–1170 (2020).
2. Gilbert, N.: A simulation of the structure of academic science. Sociol. Res. Online **2**(2), 91–105 (1997).
3. Bettencourt, L.M.A., Kaiser, D.I.: Formation of scientific fields as a universal topological transition. arXiv preprint. arXiv:1504.00319 (2015).
4. De Langhe, R.: An agent-based model of Thomas Kuhn's The Structure of Scientific Revolutions. Hist. Soc. Res. **43**(1), 28–47 (2018).
5. Gilbert, N., Ahrweiler, P., Pyka, A.: Simulating knowledge dynamics in innovation networks. Springer, Berlin (2014).
6. Brummitt, C.D., Chatterjee, S., Dey, P. S., Sivakoff, D.: Jigsaw percolation: What social networks can collaboratively solve a puzzle? Ann. Appl. Probab. **25**(4), 2013–2038 (2015).
7. Fink, T. M. A., Reeves, M., Palma, R., Farr, R. S.: Serendipity and strategy in rapid innovation. Nat. Commun. **8**(1), 2002 (2017).
8. Solé, R. V., Corominas-Murtra, B., Fortuny, J.: Diversity, competition, extinction: The ecophysics of language change. J. R. Soc. Interface **7**(53), 1647–1664 (2010).
9. Schlaile, M. P., Zeman, J., Mueller, M.: It's a match! Simulating compatibility-based learning in a network of networks. J. Evol. Econ. **28**(5), 1111-1150 (2018).
10. Jia, T., Wang, D., Szymanski, B. K.: Quantifying patterns of research-interest evolution. Nat. Hum. Behav. **1**(4), 0078 (2017).
11. Deffuant, G., Amblard, F., Weisbuch, G. Faure, T.: How can extremism prevail? A study based on the relative agreement interaction model. JASSS **5**(4), (2002).
12. National Academy of Sciences. Enhancing the postdoctoral experience for scientists and engineers. The National Academies Press, Washington (2000).
13. Kolodny, O., Creanza, N., Feldman, M. W.: Game-changing innovations: How culture can change the parameter of its own evolution and induce abrupt cultural shifts. PLoS Comput. Biol. **12**(2), e1005302 (2016).
14. Castellano, C., Fortunato, S., Loreto, V.: Statistical Physics of Social Dynamics. RMP **81**(2), 591–646 (2009).
15. Smolla, M., Akçay, E.: Cultural selection shapes network structure. Sci. Adv. **5**(8), eaaw0609 (2019).
16. Barabasi, A. L., Jeong, H., Neda, Z., Ravasz, E., Schubert, A., Vicsek, T.: Evolution of the social network of scientific collaborations. Physica A **311**(3-4), 590–614 (2002).
17. Levinthal, D. A.: Adaptation on rugged landscapes. J. Manag. Sci. **43**(7), 934–950 (1997).
18. Conrad, N., Helfmann, L., Zonker, J., Winkelmann, S., Schütte, C.:. Human mobility and innovation spreading in ancient times: A stochastic agent-based simulation approach. EPJ Data Science **7**(1), 1–22 (2018).
19. James, C., Pappalardo, L., Sirbu, A., Simini, F.: Prediction of next career moves from scientific profiles. arXiv preprint. arXiv:1802.04830 (2018).
20. Pappalardo, L., Simini, F., Rinzivillo, S., Pedreschi, D., Giannotti, F., Barabasi, A.-L. Returners and explorers dichotomy in human mobility. Nat. Commun. **6**(1), 8166 (2015).
21. Battiston, F., Cencetti, G., Iacopini, I., Latora, V., Lucas, M., Patania, A., Young, J.-G., Petri, G.: Networks beyond pairwise interactions: Structure and dynamics. Phys. Rep. **874**, 1-–92 (2020).

22. Galam, S.: Sociophysics: A physicist's modelling of psycho-political phenomena. Springer, New York (2012).
23. Sobkowicz, P.: Quantitative agent based model of user behavior in an internet discussion forum. PLoS ONE **8**(12), e80524 (2013).
24. Brandt, A. K., Eagleman, D.: The runaway species: How human creativity remakes the world. Canongate Books Ltd, Edinburgh.
25. Korhonen, J., Kasmire, J.: Adder: A simplified model for simulating the search for innovations. In: DRUID Academy Conference 2013, Aalborg, Denmark (2013).
26. Yang, V. C., Galesic, M., McGuinness, H., Harutyunyan, A.: Dynamical system model prodicts when social learners impair collective performance. PNAS **118**(35), e2106292118 (2021).
27. Centola, D., Becker, J., Brackbill, D., Baronchelli, A.: Experimental evidence for tipping points in social convention. Science **360**(6393), 1116–1119 (2018)
28. Brughmans, T., Peeples, M.: Trends in archaeological network research: A bibliometric analysis. J. Hist. Netw. Res. **1**, 1–24 (2017).
29. Derex, M., Perreault, C., Boyd, R.: Divide and conquer: Intermediate levels of population fragmentation maximize cultural accumulation. Phil. Trans. R. Soc. B **373**(1743), 20170062 (2018).
30. Kao, A. B., Couzin, I. D.: Modular structure within groups causes information loss but can improve decision accuracy. Phil. Trans. R. Soc. B **374**(1774), 20180378 (2019).