A Novel Model for Novelty: Modeling the Emergence of Innovation from Cumulative Culture

Natalie Kastel[1,2,3] and Guillaume Dumas[3,4](✉)

[1] Amsterdam Brain and Cognition Centre, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
Natalie.Kastel@PPSP.team

[2] Institute for Advanced Study, University of Amsterdam, Oude Turfmarkt 147, 1012 GC Amsterdam, The Netherlands

[3] Precision Psychiatry and Social Physiology Laboratory, CHU Sainte-Justine Research Center, Department of Psychiatry, Université de Montréal, Montréal, QC, Canada
Guillaume.Dumas@PPSP.team

[4] Mila - Quebec AI Institute, Université de Montréal, Montréal, QC, Canada

Abstract. While the underlying dynamics of active inference communication and cumulative culture have already been formalized, the emergence of novel cultural information from these dynamics has not yet been understood. In this paper, we apply an active Inference framework, informed by genetic speciation, to the emergence of innovation from a population of communicating agents in a cumulative culture. Our model is premised on the idea that innovation emerges from accumulated cultural information when a collective group of agents agree on the legitimacy of an alternative belief to the existing (or- status quo) belief.

Keywords: Active inference · Innovation · Communication · Cumulative culture · Cultural dynamics

1 Introduction

The dynamics underlying cultural evolution include the introduction of novel cultural information to a population (i.e., innovation), the transmission of established cultural information within a population (i.e., communication), and its change in prevalence (i.e., cumulative culture) (Kashima et al. 2019).

While there is a fast growing body of theoretical and empirical work on characterizing these dynamics (Aunger 2001; Buskell et al. 2019; Bettencourt et al. 2006; Creanza et al. 2017; Dawkins 1993; Dean et al. 2014; Dunstone and Caldwell 2018; Enquist et al. 2011; Gabora 1995; Heylighen and Chielens 2009; Kashima et al. 2019; Richerson et al. 2010; Stout and Hecht 2017; Weisbuch et al. 2009) mathematical models able to integrate this data into quantifiable models are scarce.

In 2015, Friston & Frith provide a quantitative model of joint communication and show that communication couples the internal states of active-inference agents and

C. L. Buckley et al. (Eds.): IWAI 2022, CCIS 1721, pp. 222–234, 2023.
https://doi.org/10.1007/978-3-031-28719-0_16

underwrites a minimal form of generalized synchrony between their internal states at a level of abstraction that allows us to characterize a statistical coupling even for agents operating with fundamentally different underlying neurobiological structures. Kastel and Hesp (2021) build on this and cast cultural transmission as a bi-directional process of communication. The idea is that when active inference agents communicate, they are able to understand each other by referring to their own generative model and inferring the internal state of the other from their behavior. This couples communicating agents in an action perception cycle of prediction and inference that induces a generalized synchrony between their internal states. Kastel & Hesp operationalise this generalized synchrony as a particular convergence between the internal states of interlocutors such that distinct belief states converge into one shared belief, and in that sense modify both of the original “parent” belief states.

The simulation of these local communication dynamics (and specifically-the convergence and subsequent modification of the belief state of each communicating agent) also serves as the basis from which to build a full blown cumulative culture model (Kastel and Hesp 2021). Cumulative culture is an emerging and prominent theory of cultural evolution which describes cultural traits as being slightly modified with every transmission such that over time these modifications accumulate to bring about an adaptive culture. Though cumulative culture is a powerful theory in that it faithfully represents the complex nature of societal change, this complexity is exceptionally challenging to formalize in quantitative models. Kastel & Hesp provided an active-inference formalization of cumulative culture by casting it as the emergence of accumulated modifications to cultural beliefs from the local efforts of agents to converge on a shared narrative. As a proof of principle for this hypothesis, they simulate a population of agents that interchangeably engage in dialogue with each other over time. When a divergent belief state is introduced to a uniform population holding (variations of) a status quo belief, it spreads through it and brings about a cumulative collective behavior of separation and isolation between groups holding distinct beliefs.

While they provide a sufficient formulation of the way slight modifications to cultural information occur during communication (previously understood as transmission) and shown how the accumulation of these dynamics affect an entire population (i.e., cumulative culture), Kastel & Hesp did not provide an account of the way novel information (i.e., the hypothesized belief state) is introduced into a population to begin with.

Within a cumulative culture framework, innovation is interpreted as the emerging property of a complexity of exchanges between agents, as opposed to the result of the mental effort of an exceptionally skillful individual. Indeed, emerging theories have put forward the suggestion that inventors and entrepreneurs are not “the brains” behind a creative idea, but are the product of a collective cultural brain (Muthukrishna and Henrich 2016). Their ideas do not stand in competition or comparison with other agents in the population, but are better understood as a nexus for previously isolated ideas within it. This cumulative approach to cultural innovation is supported by empirical findings showing that innovation rates are higher in cultures with high sociality (i.e. large and highly interconnected populations that offer exposure to more ideas), transmission fidelity (i.e. better learning between agents) and transmission variance (i.e. a willingness to somewhat deviate from the accepted learned norms (Muthukrishna and Henrich 2016).

The theory of innovation as the emerging property of a complex cultural "brain" is compatible with the theory of cumulative culture and with empirical data, but it does not provide a specific account of the mechanism by which innovation may be achieved and novel cultural beliefs and practices introduced into a population.

This paper provides an active-inference based theoretical account of Innovation as the emergence of novel cultural information from a cumulative culture. This novel account of innovation derives inspiration from the way novelty emerges in biology, namely, through a process of genetic speciation.

2 The Emergence of Innovation

2.1 Speciation in Biology

We propose that it may prove constructive to draw on specific analogies between the way novel cultural beliefs and practices emerge within a culture and the emergence of a new species in the context of biological evolution, while remaining sensitive to points at which such analogies break down. In nature, speciation occurs when a group of organisms from a particular species are separated from their original population, thus encouraging the development of their own unique characteristics. These new characteristics increasingly differentiate the two population groups when their differences grow larger as the two groups reproduce separately due to their separate environments or characteristics. Across generations, genetic differences between the old and new group become so large that they are no longer able to create offspring (i.e. mechanisms of reproductive isolation), thus highlighting the status of the subgroup as an entirely new species in its own right (Rundle and Nosil 2005).

A classic example of speciation is that of the Galápagos finch. Different species of this bird inhabit separate environments, located on different islands of the Galápagos peninsula. Over time and numerous generations, separate populations of finches developed a variety of beak morphologies, each group's morphology appearing to be adapted specifically to the feeding opportunities available on their island. While one group had developed long and thin beaks, ideal for probing cactus flowers without getting injured by the cactus, other finches developed large and blunt beaks that were perfect for nut cracking. Due to the reproductive isolation of these birds (geographic based, in this case), they developed into separate species with their own unique features.

2.2 Innovation as Cultural Speciation

Before discussing possible similarities between biological speciation and cultural innovation, a crucial difference between them should be noted. Mechanisms of biological reproductive isolation prevent members of different species from producing offspring or, in edge cases, render such offspring sterile (Palumbi 1994). In contrast, cultural evolution frequently involves cross-talk between different branches of the cultural tree – as different cultures have tended to co-opt and refurbish each other's beliefs and practices. While horizontal gene transfer is exceedingly rare in biology, cultural evolution has experienced some of its greatest accelerations precisely due to transmission of beliefs and practices across diverse cultures.

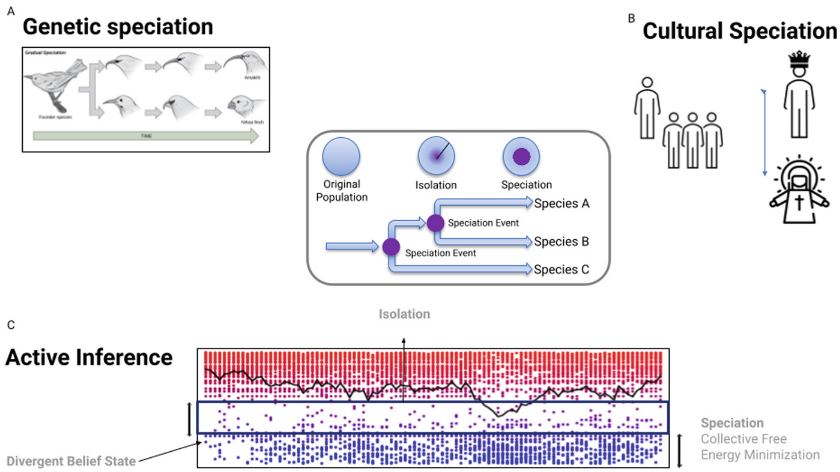

Fig. 1. A visual representation of speciation in genetics and its model in culture and active inference. Each model of speciation requires the existence of an original population that diverges through a process of group isolation such that each group develops its own unique characteristics and features. (A) genetic speciation in the galapagos finch. Geographic reproductive isolation broke up the original population of finches into those inhabiting separate islands with different selective pressures. Due to the differences in selective pressures, the separated populations developed a variety of beak morphologies that distinguished them from each other. (B) A model of speciation in religious practices. The divergence of early christians from the established jewish religion on the basis of differing interpretations of jewish scripture is modeled as a form of cultural isolation. While early Christians interpreted jewish eschatology as foretelling the arrival of a divine jewish seviour, conservative jews interpreted the same scripture as ascribing royalty to this envisioned liberator, but not divinity of any kind. As these separate streams of cultural beliefs and practices developed their own unique set of characteristics (i.e., traditions, beliefs and followers) they were no longer recognisable as part of the same religion at all and Christianity emerged as an established religion. (C) An active inference model of speciation in a cumulative culture (Kastel and Hesp 2021). When an intractably divergent belief state is introduced to a largely status quo population, locally parameterised efforts to minimize free energy bring about a self organized divergence in the population, which aligns with the process of reproductive isolation. Speciation is qualitatively observed in these simulations and is plausible under an active inference framework when representations within belief groups homogenize (i.e. shared expectations between agents emerge from a collective effort to minimize free energy). For detailed information on the methodology, and architecture of the generative model used to generate these simulations see Appendix A & B).

This crucial difference may be seen as a threat to a possible analogy between speciation and innovation because the former is made possible by virtue of a complete isolation between subgroups of a population, while such rigid isolation is not usually the case in culture. In theory this might lead to the logical conclusion that cultural speciation is simply not a possibility, because different branches of the cultural tree would not be able to maintain their characteristic integrity (as different species do) when external influences are so prevalent that they threaten any possibility for group level cultural stability.

Despite this undoubtedly logical concern, we have hard indisputable observational evidence of the existence of different nations, religions and cultural practices that have maintained their integrity for hundreds and even thousands of years in spite of cultural cross talk. This teaches us that novelty in culture is able to emerge despite perturbations from external forces to cultural 'bubbles' of communicative isolation.

For instance, Judaism is an example of a religion that has rather successfully maintained its cultural integrity despite being highly susceptible to external cultural and religious influences throughout history. That being said, practicing Judaism today is still unmeasurably different than it would have been 3000 years ago, a fact pointing to internal changes in culture while its characterizing features and socio-cultural boundaries remain at least partially intact. This can be well explained by an account of cumulative cultural dynamics (Kastel and Hesp 2021). Internal changes to the Jewish religion can be attributed to an accumulation of incremental modifications that occurred with every generation (and within generations) through the communication of religious practices within the community, with relatively minimal (though still unavoidable) blending and mixing with Roman or Christian religious traditions and beliefs.

According to our account so far, each transmitted belief is translated and fitted to a specific phenotype-congruent representation of that belief state on the receiving end of the exchange (Kashima et al. 2019). Individual "private" representations of cultural beliefs therefore fuse together to create new representations of old traditions, such that cultural reproduction consists in the facilitation of different subjective representations of the same belief state, where "sameness" is in turn derived completely from subjective interpretations inherent to each individual's communicative capacity (e.g., their language). Our account is therefore capable of capturing the way a culture evolves internally, without much need for (and even despite!) external influences.

Importantly, we suggest that it might also be possible to describe and even model the way innovation emerges from these dynamics, in a form of cultural speciation.

Our model leaves the notion of a cultural speciation purposefully abstract, Such that it may take on the form of any event, fashion, ideology, preference, language or behavior that, in time, separates between two identifiable streams of culture.

To give a concrete example of what might be meant by this, we return to the slightly thorny subject of religion, and specifically, the speciation (i.e., cultural differentiation) of Christianity in the context of Judaism. This discussion is intended as an illustration of divergence of belief and practices, without judging the value of their content per se, in light of a specific point of disagreement. It should not be taken as an exhaustive treatment of differences between Judaism and Christianity. Historically, early Christians diverged from the established Jewish religion, at least partly on the basis of differing interpretations of Jewish scripture (as referred to in panel (B) of Fig. 1). According to Jewish eschatology (i.e., Jewish scholars' interpretations of their scriptures with regard to the end of times) At the time, a Jewish king referred to as "messiah" (savior) and "son of David" (a descendent from the Davidic line) would rise to rule the Jewish people and bring them salvation from their hardships. Early Christians believed that Jesus of Nazareth fulfilled the criteria for being that promised savior (Lucass 2011, p. 4–6) A critical divergence between the interpretation of early Christians and conservative Jews, was that the characteristics of the "messiah" – according to the interpretation of conservative Jews – did not (and for religious reasons, could not) include divinity of any kind but merely plenty of charisma and leadership that would lead to royalty. The emergence of these distinct interpretations, or incompatible representations for the belief in the prophecy of a messiah could be seen as the speciation event that accelerated the differentiation of Christianity from Judaism (Fig. 1B). As one example of direct behavioral incompatibility stemming from these divergent interpretations, early Christians believed that Jesus had "fulfilled" the religious law brought by Moses and quickly discarded adherence to, e.g., Jewish dietary laws, circumcision, and sacrificial practices (see, e.g., the Council of Jerusalem described in the Book of Acts, chapter 15; estimated to have taken place around 50 AD). Cultural reproductive isolation followed as these separate streams of cultural beliefs and practices developed their own unique set of characteristics, traditions, beliefs and followers – until they were no longer recognisable as part of the same religion at all and Christianity emerged as an established religion (Boyarin 1999, p. 1–22). An interesting recent example of further "speciation event" in this regard is the recent (20th century) emergence of "Messianic Judaism", a syncretic Christian movement that mixes adherence to Judaic laws with acceptance of Jesus Christ as the messiah (Melton 2005).

This is only one example of the type of speciation that might take place within a cultural arena, and it represents a direct analogy to one of four types of biological speciation (Rundle and Nosil 2005). Our particular example corresponds to sympatric speciation, which occurs when genetic polymorphism causes two groups from the same species to evolve differently until they can no longer interbreed and are considered separate species. Our Galápagos finch example, on the other hand, was an example of allopatric biological speciation, in which a particular geographic barrier prevents groups of the same species from interbreeding until they undergo genotypic divergence to the point of reproductive isolation, where they are also considered different species. Our simulations focus on the sympatric form of cultural speciation, although we assume that parallels can be made for all four types of biological speciation.

2.3 Innovation in Active Inference

In active inference, the attunement of interlocutor's generative models on the microscale translates over time and with multiple encounters into collective free energy minimisation on the macroscale. Kastel and Hesp (2021) (Fig. 1C) show that simulations of cumulative culture are aligned with this premise when locally parameterised efforts to minimize free energy by individual agents bring about a self organized separation in the population when an intractably status-quo-divergent belief is introduced. In other words, it would appear that simulations cumulative cultural dynamics imitate reproductive isolation when separate belief groups (i.e. blue and red in these simulations) diverge and communicate in observable isolation from one another. Cultural speciation, while not specifically observed in these simulations, is plausible under an active inference framework when representations within belief groups homogenize (i.e. shared expectations between agents emerge).

While this paper is limited to theorizing about the emergence of innovation from a cumulative culture diffusion, there is a great deal of potential for future modeling work in this field. Such work should include at least two added components to our theoretical account in order to provide a consistent and complete model of the emergence of innovation from the cumulative culture dynamics provided by Kastel and Hesp (2021).

First, both reproductive (i.e., communicative) isolation and speciation need to be simulated in a formalized model of active inference. Communicative isolation between groups holding divergent beliefs should be quantifiable and should emerge naturally from local (agent based) free energy minimization. Similarly, cultural speciation should be operationalised and simulated as collective free energy minimization that brings about the emergence of shared representations within groups.

Secondly, and most importantly, A novel belief state should naturally emerge from these dynamics as opposed to being synthetically introduced into the population, as was done in our simulations. This should involve the design of a new paradigm for simulating the natural emergence of a sufficiently divergent belief state from the dynamics of cumulative culture, namely, from the accumulation of modifications to cultural information. What this means is that contrary to being mechnichally placed in the population in a manner that allows belief states to remain abstract, a belief state that emerges from several modifications on it will need to refer back to the complex content (namely, the manifold of alterations) that brought it about.

For the emerging belief state to be "sufficiently divergent" from a status quo population, it needs to be dissimilar to the status quo belief to a degree that is large enough that it creates a desired level of isolation between belief-groups (i.e. each group maintains its integrity), but not so large that communication with the first agent holding this belief becomes completely impossible. Note that the latter condition is simply the assumption in active inference for the possibility of communication. Namely, that for agents to arrive at a hermeneutic resolution and be able to understand each other, they must employ sufficiently similar generative models (Friston and Frith 2015). When this is not the case, and agents employ intractably dissimilar cultural beliefs, they will not be able to refer to their own generative model to infer the internal state of another from their behavior.

We arrive at an interesting, Goldilocks precondition for the emergence of a novel belief from cumulative cultural dynamics. A novel belief that emerges from the mixing and merging of beliefs over time should be neither too divergent, nor too similar to the status quo belief. When the former is the case, the agent that suddenly emerges with a potentially novel belief, has in his mind an idea so unique and exclusive that it is incomprehensible and unrelatable to other members of the population. On the other hand, if the latter is the case, and an agent emerges with a belief that is too similar to the status quo, his belief does not differ enough from the status quo to be isolated from it as a separate stream.

In conclusion, the theory of innovation we have discussed defines exactly the difference between a belief state that is only slightly modified during communication, and one that is considered novel. For cultural beliefs and practices to be slightly modified such that they continue to evolve, they need only comply with the hermeneutic condition and allow for communication between agents carrying this information to exchange ideas. This conclusion is derived from the theories and formulations of communication and cumulative culture brought forward in Kastel and Hesp (2021). Innovation, however, seems to have harsher requirements. It needs to comply with both the hermeneutic condition (i.e., needs to be sufficiently similar to the status quo) as well as an "isolation condition" (i.e., needs to be sufficiently divergent from the status quo). This conclusion is derived from theories brought forward in this paper, which provides a solid foundation on which to build a complete account of cultural dynamics.

A formalized account of innovation as an emergent property from the cumulative dynamics presented in this proposal would bring the circular dynamics of a complex culture to a satisfying close. Under such an account, not only would cumulative culture naturally emerge from a complex communication network of agents (as shown in Kastel and Hesp (2021)), but innovation would emerge from cumulative culture and underlie communication in a repeating, recursive loop that is the hallmark of complex dynamical systems.

3 Conclusions

We discuss a possible theory of innovation as the emergent property of cumulative cultural dynamics. We suggest that innovation emerges when gradual modifications to cultural information spontaneously produce a "sufficiently divergent" belief state that meets a goldilocks condition of being neither too similar, nor too conflicting with the status quo in the population. If the former is not met, communication with the agent holding this belief will not result in coordination with members of the existing population, and the alternative belief will not propagate. If the latter is not met, its propagation will not create the level of isolation necessary between belief groups for each group to maintain its integrity to be considered novel at all.

Appendix A - Methodology for Simulating the Dynamics of Cumulative Culture

A.1 Simulating the Local Dynamics of Communication

In our model, cultural transmission is cast as the mutual attunement of actively inferring agents to each other's internal belief states. This builds on a recent formalization of communication as active inference (Friston and Frith 2015) which resolves the problem of hermeneutics, (i.e., provides a model for the way in which people are able to understand each other rather precisely despite lacking direct access to each other's internal representations of meaning) by appealing to the notion of generalized synchrony as signaling the emergence of a shared narrative to which both interlocutors refer to. In active inference, this shared narrative is attained through the minimisation of uncertainty, or (variational) free energy when both communicating parties employ sufficiently similar generative models. We build on this to suggest that having sufficiently similar generative models allows communicating agents to recombine distinct representations of a belief (expressed as generative models) into one synchronized, shared model of the world (Fig. 2). When we simulate the belief-updating dynamics between interacting agents, the cultural reproduction of a particular idea takes the form of a specific convergence between their respective generative models.

Under this theory, the elementary unit of heritable information takes the form of an internal belief state, held by an agent with a certain probability. When we simulate the belief-updating dynamics between interacting agents, a reproduced cultural belief is carried by the minds (or generative models) of both interlocutors as a site of cultural selection, where it may be further reproduced through the same process. Our simulations of communication involve two active inference agents with distinct generative models and belief claims that engage in communication over a hundred time steps.

A.2 Simulating the Global Dynamics of Cumulative Culture

Cultural beliefs and practices spread within a society through communication, a process which we have referred to as the local dynamics of cumulative culture. This description is appropriate because the accumulated outcomes of each (local) dyadic interaction collectively determine the degree to which an idea is prevalent in a culture. Moving from local communication dynamics to a degree to which an idea is prevalent in a cumulative culture is what we refer to as the global dynamics of cumulative culture.

In our simulations of a cumulative culture, 50 active inference agents simultaneously engage in local dyadic communication as shown in our first simulation, such that 25 couples are engaged in conversation at every given time step. At the first time step, all agents have relatively similar belief states- referred to as the status quo. When we introduce an agent holding a divergent belief state to that of the status quo in the population, it propagates through it via pseudo-random engagements of agents in dialogue. In a simulated world of actively inferring agents, their individual mental (generative) models are slightly modified with every interlocutor they encounter, as their distinct representations converge to a shared narrative (Constant et al. 2019). The attunement of interlocutor's to each other's generative models on the microscale thus translates over time and with multiple encounters into collective free energy minimisation on the macroscale.

Appendix B - Generative Model Architecture, Factors and Parameters

In our simulations, agents attempt to convince each other of a cultural belief by utilizing generative models that operate with local information only. For the establishment of such generative models, we will formulate a partially observed Markov decision process (MDP), where beliefs take the form of discrete probability distributions (for more details on the technical basis for MDP'S under an active inference framework, see Hesp 2019).

Under the formalism of a partially observed Markov decision process, active inference entails a particular structure. Typically, variables such as agent's hidden states (x, s), observable outcomes (o) and action policies (u) are defined, alongside parameters (representing matrices of categorical probability distributions).

B.1 Perceptual Inference

The first level of this generative model aims to capture how agents process belief claims they are introduced to through conversation with other agents. The perception of others' beliefs (regarded in active inference as evidence) requires prior beliefs(represented as likelihood mapping A1 about how hidden states (s1) generate sensory outcomes (o).

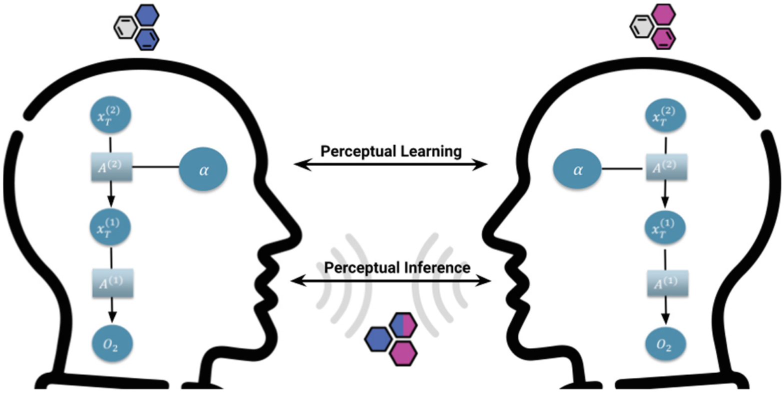

Fig. 2. Communication Coupling Parameters. Our model defines two groups of parameters that couple the internal states of agents: Learning and inference. Perceptual learning (A2) is the learning of associations between emotional valence and belief states that guide the long term actions of our agents who hold and express beliefs. This learning happens at slow time scales, accumulating across multiple interactions and used to modify models over extended periods of exchange. Perceptual Inference (A1) – namely, sensitivity to model evidence – operates on fast time scales and is direct and explicit to agents during dialogue. Importantly, we hypothesized that without precise evidence accumulation, agents would be insensitive to evidence regarding the belief state of the other, and their internal states would not converge.

Specifically, our agents predict the likelihood of perceiving evidence toward a particular expressed belief, given that this belief is "the actual state of the world". Parameterizing an agent's perception of an interlocutor's expression of belief in terms of precision values can be simply understood as variability in agents' general sensitivity to model evidence. High precisions here correspond to high responsiveness to evidence for a hidden state and low precisions to low responsiveness to evidence. Precisions for each agent were generated from a continuous gamma distribution which is skewed in favor of high sensitivity to evidence on a population level (See Fig. 2 & Fig. 3: Perception).

B.2 Anticipation

At this level, our generative model specifies agents' beliefs about how hidden states (detailed in Appendix A2) evolve over time. State transition probabilities [B1] define a particular value for the volatility of an agent's meeting selection (s2) and belief expression (s1) [B1]. For each agent, this precision parameter is sampled from a gamma distribution, determining the a priori probability of changing state, relative to maintaining a current state. Note that belief states themselves are defined on the continuous range $<0, 1>$ (i.e., as a probability distribution on a binary state), such that multiplication tends to result in a continuous decay of confidence over time in the absence of new evidence (where the rate of decay is inversely proportional to the precision on B) (See Fig. 3: Anticipation).

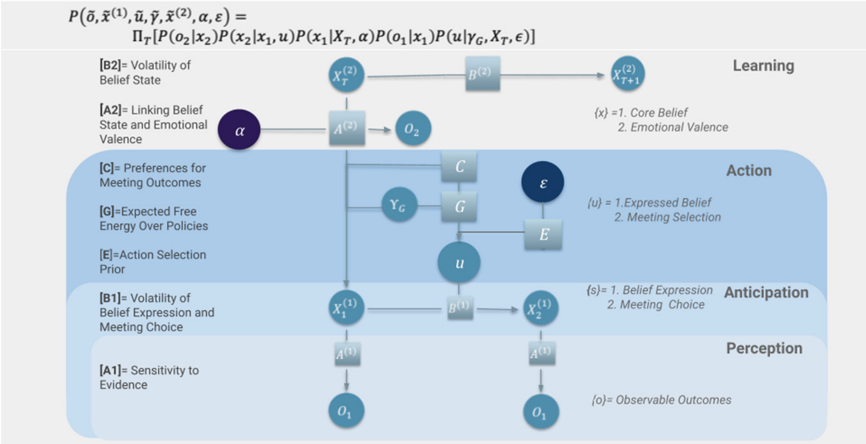

Fig. 3. A generative model of communication. Variables are visualized as circles, parameters as squares and concentration parameters as dark blue circles. Visualized on a horizontal line from left to right-states evolve in time. Visualized on a vertical line from bottom to top- parameters build to a hierarchical structure that is in alignment with cognitive functions. Parameters are described to the left of the generative model and variables are described on the right.

B.3 Action

After perceiving and anticipating hidden belief states in the world, our agents carry out deliberate actions biased towards the minimum of the expected free energy given each action (a lower level generative model for action is detailed in Appendix A4 and A5). At each time point, a policy (U) is chosen out of a set of possible sequences for action. In our simulations, two types of actions are allowed: selecting an agent to meet at each given time point (u2) and selecting a specific belief to express in conversation (u1). The first allowable action holds 50 possible outcomes (one for each agent in the simulation) while the second is expressed on the range <0, 1>, where the extremes correspond to complete confidence in denying or supporting the belief claim, respectively. Each policy under the G matrix specifies a particular combination of action outcomes weighted by its expected negative free energy value and a free energy minimizing policy is chosen (See Fig. 3: Action).

Voluntary Meeting Selection. While the choice of interlocutor is predetermined in a dyad, our multi-agent simulations required some sophistication in formulating the underlying process behind agents' selection for a conversational partner (s3) at each of the hundred time points. Building on previous work on active inference navigation and planning (Kaplan and Friston 2018), agents' meeting selection in our model is represented as a preferred location on a grid, where each cell on the grid represents a possible agent to meet.

We demonstrate the feasibility of incorporating empirical cultural data within an active inference model by incorporating (1) confirmation bias through state-dependent preferences [C], biasing meeting selection through the risk component of expected free energy (G) and (2) novelty seeking through the ambiguity component of expected free energy. The first form of bias reflects the widely observed phenomenon in psychology research that people's choices tend to be biased towards confirming their current beliefs (Nickerson 1998). The second form of bias reflects the extent to which agents are driven by the minimisation of ambiguity about the beliefs of other agents, driving them towards seeking out agents they have not met yet. For a detailed account on the process of meeting selection in these simulations, the reader is referred to Kastel and Hesp 2021.

B.4 Perceptual Learning

On this level agents anticipate how core belief states (specified in Appendix A1) might change over time [B2] (Fig. 2.3). This is the highest level of cognitive control, where agents experience learning as a high cognitive function (higher level generative model is detailed in Appendix A3). By talking with other simulated agents and observing their emotional and belief states, our agents learn associations between EV and beliefs via a high level likelihood mapping [A2], (updated via concentration parameter α). The Updating of core belief, based on beliefs expressed by other agents, is detailed in Appendix A7. This learning is important because it provides our agents with certainty regarding the emotional value they can expect from holding the alternative belief to the status quo, which has low precision at the beginning of the simulation (before the population is introduced to an agent proclaiming this belief). The prior P(A) for this likelihood mapping is specified in terms of a Dirichlet distribution.

References

Aunger, R.: Darwinizing culture: the status of memetics as a science (2001)

Buskell, A., Enquist, M., Jansson, F.: A systems approach to cultural evolution. Palgrave Commun. **5**(1), 1–15 (2019)

Bettencourt, L.M., Cintron-Arias, A., Kaiser, D.I., Castillo- Chávez, C. The power of a good idea: quantitative modeling of the spread of ideas from epidemiological models. Phys. A: Stat. Mech. Appl. **364**, 513–536 (2006)

Boyarin, D.: Dying for God: Martyrdom and the Making of Christianity and Judaism. Stanford University Press (1999)

Creanza, N., Kolodny, O., Feldman, M.W.: Cultural evolutionary theory: how culture evolves and why it 562 matters. Proc. Natl. Acad. Sci. **114**(30), 7782–7789 (2017)

Dawkins, R.: Viruses of the mind. Dennett Crit. Demystifying Mind **13**, e27 (1993)

Dunstone, J., Caldwell, C.A.: Cumulative culture and explicit metacognition: a review of theories, evidence and key predictions. Palgrave Commun. **4**(1), 1–11 (2018)

Enquist, M., Ghirlanda, S., Eriksson, K.: Modeling the evolution and diversity of cumulative culture. Philos. Trans. Roy. Soc. B: Biol. Sci. **366**(1563), 412–423 (2011)

Friston, K., Frith, C.: A duet for one. Conscious. Cogn. **36**, 390–405 (2015)

Gabora, L.: Meme and variations: A computational model of cultural evolution. In: 1993 Lectures in Complex Systems, pp. 471–485. Addison Wesley (1995)

Gordon Melton, J.: Encyclopedia of Protestantism. Infobase Publishing (2005)

Heylighen, F., Chielens, K.: Evolution of culture, memetics. In: Meyers, R. (ed.) Encyclopedia of Complexity and Systems Science, pp. 3205–3220. Springer, New York (2009). https://doi.org/10.1007/978-0-387-30440-3_189

Kashima, Y., Bain, P.G., Perfors, A.: The psychology of cultural dynamics: What is it, what do we know, and what is yet to be known? Annu. Rev. Psychol. **70**, 499–529 (2019)

Kastel, N., Hesp, C.: Ideas worth spreading: a free energy proposal for cumulative cultural dynamics. In: Kamp, M., et al. (eds.) ECML PKDD 2021. Communications in Computer and Information Science, vol. 1524, pp. 784–798. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93736-2_55

Muthukrishna, M., Henrich, J.: Innovation in the collective brain. Philos. Trans. Roy. Soc. B: Biol. Sci. **371**(1690), 20150192 (2016)

Palumbi, S.R.: Genetic divergence, reproductive isolation, and marine speciation. Annu. Rev. Ecol. Syst. 547–572 (1994)

Richerson, P.J., Boyd, R., Henrich, J.: Gene-culture coevolution in the age of genomics. Proc. Natl. Acad. Sci. (Supplement 2), 8985–8992 (2010)

Rundle, H.D., Nosil, P.: Ecological speciation. Ecol. Lett. **8**(3), 336–352 (2005)

Stout, D., Hecht, E.E.: Evolutionary neuroscience of cumulative culture. Proc. Natl. Acad. Sci. **114**(30), 7861–7868 (2017)

Weisbuch, M., Pauker, K., Ambady, N.: The subtle transmission of race bias via televised nonverbal behavior. Science **326**(5960), 1711–1714 (2009)