Title

Multiple heads outsmart one: A computational modelfor distributed decision making

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Kapourchali, Masoumeh Heidari
Banerjee, Bonny

Publication Date

2018

Multiple heads outsmart one: A computational model for distributed decision making

Masoumeh Heidari Kapourchali (mhdrkprc@memphis.edu), Bonny Banerjee (bbnerjee@memphis.edu)
Institute for Intelligent Systems, and Department of Electrical & Computer Engineering
The University of Memphis, Memphis, TN 38152, USA

Abstract

Distributed cognition and decision making has been a topic of intense research in the recent years. In this paper, a computational model of distributed decision making using a community of predictive coding agents is developed. The agents are embodied multimodal entities and situated in a shared environment. They have different visibility of the environment due to unique sensory and generative models. We show that communication between agents helps each of them reach a shared decision in a way that cannot be reached by brain processes in a single agent. Using a simulated environment, we show that sensory limitations may lead to incorrect or delayed causal inferences giving rise to conflicts in the mind of a predictive coding agent, and communication helps to resolve such conflicts and overcome the limitations.

Keywords: distributed cognition, predictive coding, agent, embodiment, communication, free energy, active inference

1. Introduction

Distributed decision making has been a topic of intense research in quantitative decision analysis (Schneeweiss, 2012) and software agent technologies. Our work explores distributed decision-making using a community of *predictive coding* agents, each equipped with an independent and unique sensory-motor system relentlessly executing the perception-action loop. Similar to biological entities, none of the agents can completely observe the reality. Each agent's version of the world is a function of its sensory system and internal model. In the context of decision-making, each agent's limitations can be overcome to some extent through communication with the other agents. Communication extends an agent's perceptual field and allows efficient causal knowledge acquisition by sampling other agents' internal causal models.

Most of the computational models of communication can be categorized into one of the following classes: a) predefined communication protocols, b) planning or learning methods, c) evolution or reinforcement learning, and d) cooperative or competitive settings (Sukhbaatar, Fergus, et al., 2016; Foerster, Assael, de Freitas, & Whiteson, 2016; Foerster et al., 2016). Friston and Frith (2015a, 2015b) model the communication between two predictive coding agents using the *hermeneutic cycle*. Their goal is to facilitate learning a model of each agent by the other, assuming the generative and sensory systems of the two agents to be the same. Chen et al. (2016) used the hermeneutic cycle to model interactive behaviors between two robots. Advantages of predictive coding based on free energy minimization over traditional approaches of optimizing agent's behaviors are discussed in (Friston, Daunizeau, & Kiebel, 2009; Friston et al., 2013).

In this paper, we propose a computational model of distributed decision making among multiple predictive coding agents through mutual communication about a common subject. In Figure 1, this problem is illustrated using two agents, A_1 and A_2, communicating regarding a third entity, A_3. The sensory/generative system of each agent is unique. To make the problem interesting, each agent is assumed to be multimodal, receiving sensory observation in two modalities: one, directly from the common subject, and the other from the other agent(s) due to communication about the subject. The former modality is unique for each agent while the latter modality (for receiving communicative inputs) is common across all agents. At any instant of time, the goal is for all communicating agents to reach a decision on the state of the common subject more accurately and quickly than each of them could have by itself. This goal stems from social cognition research (Di Paolo & De Jaegher, 2012) where communication is construed as dynamic interaction among multiple individuals which helps reach a shared decision in a way that could not be reached by brain processes in a single individual. Communication in our model is at the level of agents' beliefs and is not limited to low-level brain/spinal signals.

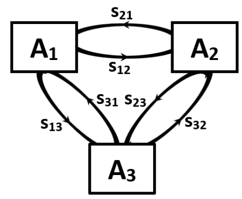

Figure 1: Distributed decision-making through communication between two agents, A_1 and A_2, regarding the state of a third entity, A_3. s_{ij} denotes a signal passed from entity A_i to entity A_j which could be a communicative message or a sample of an observable variable related to the state of an entity.

The rest of the paper is organized as follows. Section 2 introduces the necessary definitions. The proposed model is described in Section 3, including the computational model of a single agent and the model for communication between agents. The experimental results are discussed in Section 4.

2. Definitions

The terms and concepts relevant to the framework are discussed in this section.

Definition 1. (Agent) An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators (Russell & Norvig, 2002). The agent discussed in this paper is implemented in software; its actions are limited to sampling the environment (simulated) and communication with other similar agents.

Definition 2. (Hermeneutic circle) The hermeneutic circle is used as a model of explaining communication. It refers to the problem of circularity of understanding (Chandler & Munday, 2011) where understanding the first agent presupposes understanding the second agent, which in turn presupposes understanding the first agent (Friston & Frith, 2015a).

Definition 3. (Variational free energy) Variational free energy is a measure of salience based on the divergence between the recognition $q(x)$ and generative density $p(\varphi, x)$ (Friston et al., 2009): $F = - < \ln p(\varphi, x) >_q + < \ln q(x) >_q$.

Definition 4. (Recognition density) Recognition density $q(x)$, is a probabilistic representation of causes which is encoded by internal states μ. Assuming it as a Gaussian density, it is also called Laplace approximation (Friston, 2010): $q(x) = \mathcal{N}(x; \mu, \zeta) = \frac{1}{\sqrt{2\pi\zeta}} exp\{-(x-\mu)^2/2\zeta\}$

Definition 5. (Generative density) Generative density $p(\varphi, x)$ is a joint probability density relating environmental states and sensory data. It is usually specified in the form of a prior $p(x)$ and a likelihood $p(\varphi|x)$ (Buckley, Kim, McGregor, & Seth, 2017).

Definition 6. (Predictive coding) The biological brain operates as a generative prediction machine which is hierarchically organized (Rao & Ballard, 1999; Friston, 2005). It has been proposed that higher-level areas predict lower-level activity while lower areas feedforward the prediction errors, thereby removing the predictable or redundant information in the input (Rao & Ballard, 1999), referred to as predictive coding. In this paper, predictive coding will be used to denote a class of models that make inferences (predictive and causal), act and learn by minimizing variational free energy. Predictive coding can be conceptually understood as the SELP cycle (Banerjee & Dutta, 2014, 2013).

Variables used in the paper are listed in Table 1.

Table 1: Symbols and notations.

Variable	Description
φ	Sensory data
μ	Belief (expectation of cause)
ε_φ	Sensory prediction error
ε_p	Prior prediction error
σ_φ	Variance of generative density
σ_p	Variance of prior density
x	Environmental variables
x_p	Mean of prior density
a	Action
f_s	Sampling frequency

3. Models and methods

In this paper, the interaction of a group of embodied agents is modeled to infer the states of the environment in which all the agents are situated. The environment is partially-observable to each agent due to their sensory limitations.

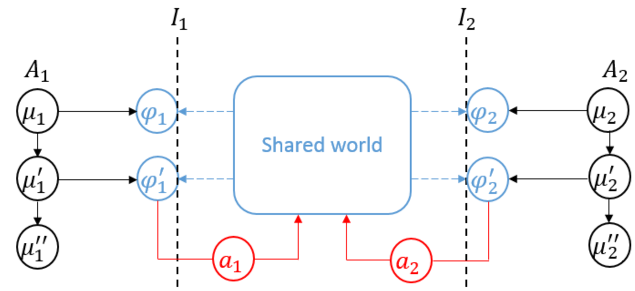

Figure 2: A schematic representation of two predictive coding agents with unique internal (generative) models in a shared environment (modified from (Buckley et al., 2017)). Interactions between generalized internal states (black) and sensory data (blue) are shown. The agents' actions on the world are represented by a_i (red). I_1, I_2 are the interfaces for agents A_1, A_2 respectively. Everything to the right of I_1, including A_2 and the shared environment, is considered as the external environment for A_1. [Best viewed in color.]

An agent and its environment

In our framework, each agent has a unique internal model and shares the environment with other agents. In addition to perception of the shared environment, each predictive coding agent is required to have a model of other agents as part of its internal model to anticipate their future actions. Each agent can act on its environment and change its state. Therefore, even though the agents are independent entities, their actions and perceptions are not entirely independent. Figure 2 shows the diagram of an environment shared by two predictive coding agents, each with a generative internal model. The two agents have unique sensors and effectors, and can act on and perceive from the shared environment. The environmental states cannot be observed directly and have to be inferred from sensory observations. Similarly, the state inferred by the other agent is also unobservable and may be estimated from the sensory observations of that agent's behaviors.

As a running example throughout this paper, consider two agents trying to infer the state of their common environment. One agent is equipped with a sensor that senses the light intensity (a.k.a. *light-agent*) while the other agent is equipped with a sensor that senses the temperature (a.k.a. *heat-agent*). At any time instant, the environment can be in one of three states: $noEvent$, $firework$ or $fire$. Each agent's goal is to infer the state of the environment at all times. The environment is modeled as:

$$f(x) = \begin{cases} noEvent, & \text{if } x < 0.1 \\ firework, & \text{if } 0.1 \leq x < 1 \\ fire, & \text{if } x \geq 1 \end{cases} \quad (1)$$

where x denotes the state of the environment. Each agent is also equipped with an actuator using which it can sample its environmental signal, such as light intensity or temperature, at a frequency of its choice within a range.

The generative density of agent A_i is given by:

$$p(\varphi_i|x) = \mathcal{N}(\varphi_i; g_i(x), \sigma_{\varphi_i}) \tag{2}$$

where $\mathcal{N}$ denotes the normal density with mean $g_i(x)$ and variance σ_{φ_i}, since the observations are noisy. The generative model g is unique for each agent; it is a mapping from the causes to the observations where the observations are function of the sensors or body of the agent. Let A_1 and A_2 be the light-agent and heat-agent respectively. Then g_1 is defined for A_1 as:

$$g_1(x) = \begin{cases} xt^{\alpha_1 - 1}(1-t)^{\beta_1 - 1}, & \text{if } x < 0.1 \\ xt^{\alpha_2 - 1}(1-t)^{\beta_2 - 1}, & \text{if } 0.1 \leq x < 1 \\ xt^{\alpha_3 - 1}(1-t)^{\beta_3 - 1}, & \text{if } x \geq 1 \end{cases} \tag{3}$$

where α_k, β_k $(k = 1,2,3)$ are predefined parameters and t denotes time. Also, g_2 is defined for A_2 following the convection equation, as:

$$g_2(x) = \begin{cases} xh(T_{hot_1} - T_{cold})Bt, & \text{if } x < 0.1 \\ xh(T_{hot_2} - T_{cold})Bt, & \text{if } 0.1 \leq x < 1 \\ xh(T_{hot_3} - T_{cold})Bt, & \text{if } x \geq 1 \end{cases} \tag{4}$$

where T denotes temperature in Kelvin, B is the area of exposure, and h is a constant. T_{hot} varies with situations such that a change in temperature due to $fire$ is different from that due to $firework$. The agents are initialized with a prior regarding the environmental states which is assumed to be a normal distribution $\mathcal{N}(\mu; v_{p_i}, \sigma_{p_i})$ with mean v_{p_i} and variance σ_{p_i} for agent A_i. It is assumed that the frequency of sampling, f_{s_i}, by A_i of its environment is proportional to the change in its observation:

$$f_{s_i} = \frac{d\varphi_i(t)}{dt} \tag{5}$$

An agent samples the environment using its body which constitutes a behavior that is observable to other agents.

Each agent can independently infer the environmental states by minimizing the free energy, given by:

$$F = \int -q(x) \ln p(x, \varphi) dx + \int q(x) \ln q(x) dx \tag{6}$$

where the first term is the average energy and the second therm is negative of entropy associated with the recognition density (Buckley et al., 2017). Assuming $q(x)$ to be a sharply peaked Gaussian density function (i.e. the Gaussian bell shape is squeezed towards a delta function), the most likely value of the environmental state is estimated iteratively using Bayesian approximation as follows:

$$\frac{\partial F}{\partial \mu} = \dot{\mu} = \varepsilon_\varphi g'(\mu) - \varepsilon_p \tag{7}$$

where ε_φ and ε_p are updated as follows:

$$\dot{\varepsilon}_p = \mu - x_p - \sigma_p \varepsilon_p \tag{8}$$

$$\dot{\varepsilon}_\varphi = \varphi - g(\mu) - \sigma_\varphi \varepsilon_\varphi \tag{9}$$

and the prediction errors are $\varepsilon_\varphi = (\varphi - g(\mu))/\sigma_\varphi$ and $\varepsilon_p = (\mu - x_p)/\sigma_p$. For a detailed derivation of equ. 7 from equ. 6, refer to (Bogacz, 2015). Note that, μ is the belief of an agent from its observation of the environment without being influenced by any agent through communication.

Reading others' minds from their behaviors

In the real world and also in our simulated environment, agents have different sets of knowledge due to differences in sensory systems/body and prior experience. Communication with other agents helps to sample from their knowledge. However, an agent may be so biased towards its own beliefs that it fails to detect its need for communication. In the context of predictive coding, it means that the agent fails to register a prediction error in which case there is no way to improve its perception.

Friston and Frith (2015a) observe that there is no way to verify whether an agent's interpreted cause of another's behavior corresponds to the latter's actual cause or not. The best the agent can do is to invent a coherent story that minimizes all conflicts in its mind. The ability to interpret an agent requires a model of that agent to be learned by observing its behaviors. In addition to predicting the environment, a predictive coding agent should be able to predict the other agents' behaviors. The observable behaviors of our light-agent and heat-agent are their sampling frequencies which are assumed to be noisy. The model of agent A_i in the mind of agent A_j is of the form: $\mu_{ij} = H(f_{s_j}; \theta_H)$ where H is a mapping from f_{s_j} to μ_{ij} given the set of parameters θ_H, and μ_{ij} is A_i's belief based on A_j's behavior which is different from that due to its observation of the shared environment.

Multiple agents and their communication

In order to extend our discussion to multiple agents, the light-agent and heat-agent will be equipped with a sensor that can sense the frequency of sampling of the environment by the other agents. Each agent has two effectors: one for sampling the environment and the other for sending communicative messages to other agents. Thus, each agent receives observations regarding the shared environment from two sources: one directly from the environment via their light/temperature sensors and the other from the communicating agent. A conflict arises in the mind of an agent whenever the inferred causes from these two sources are not in agreement. Such conflicts have to be resolved by further sampling of the environment and communication. There are many approaches in the literature for conflict resolution (Adler, Davis, Weihmayer, & Worrest, 1998; Oleary, 1999; Müller & Dieng, 2000; Sobieska-Karpińska & Hernes, 2014). We use belief revision based on trust. Trust is measured by an agent's level of confidence regarding its belief. Communication is a language that both agents ought to understand; that is, they are required to have the ability to encode and decode the communicative messages. In our running example, we assume a message to be a function of the other agent's belief (μ_j), written as $msg(\mu_j, \theta_{comm})$, where θ_{comm} is a set of parameters of the model and can be learned from data. After receiving a message from A_j, A_i's belief is revised as follows:

$$\hat{\mu_i} = \underset{x}{\operatorname{argmax}}\ p(x|\varphi_i, msg(\mu_j, \theta_{comm})) \tag{10}$$

Assuming the noise components to be independent and using Bayes rule, we get (Deneve & Pouget, 2004):

$$\begin{aligned} p(x|\varphi_i, msg) \propto p(\varphi_i, msg|x) \propto p(\varphi_i|x)p(msg|x) \\ \propto p(x|\varphi_i)p(x|msg) \end{aligned}$$

where $p(x|\varphi_i)$ and $p(x|msg)$ are Gaussian probability densities. The bimodal estimate can be a linear combination of the unimodal estimates. For N agents where all agents send messages to A_i except itself, the bimodal estimate is:

$$\hat{\mu_i} = \left(\frac{\mu_i}{\sigma_{p_i}} + \sum_{\substack{n=1 \\ n \neq i}}^{N} \frac{msg_n}{\sigma_{p_n}} \right) \Big/ \sum_{m=1}^{N} \frac{1}{\sigma_{p_m}} \tag{11}$$

Here $\hat{\mu_i}$ is the belief of A_i after communicating with other agents and weighing their messages. Inverse of σ_{p_j} is a measure of A_i's trust on A_j's message. If all weights are equal, i.e. $\sigma_{p_i} = \sigma_{p_j}\ \forall i, j, i \neq j$, the belief of all agents will converge to the same value which will render all agents except one redundant. Learning a unique model of other agents by each agent allows the entire multiagent system to store more knowledge about the shared environment and allows each agent to resolve conflicts with other agents amicably. By sampling from other agents' internal models through communication, each agent acquires causal knowledge more efficiently than by observing the environment as the environment can only present correlations but an agent can share its causal knowledge.

Inverse of variance is a measure of precision in predictive coding (Friston et al., 2009). An agent may not have an accurate model of trust from the beginning. To improve the model, the precision is updated along with minimization of free energy. The update rules for parameters of prior density with each observation are as follows (Bogacz, 2015):

$$\frac{\partial F}{\partial x_p} = \dot{x}_p = \frac{\mu - x_p}{\sigma_p} = \varepsilon_p \tag{12}$$

$$\frac{\partial F}{\partial \sigma_p} = \dot{\sigma}_p = \frac{1}{2}\left(\frac{(\mu - x_p)^2}{\sigma_p^2} - \frac{1}{\sigma_p}\right) = \frac{1}{2}(\varepsilon_p^2 - \sigma_p^{-1}) \tag{13}$$

x_p and σ_p converge to mean and standard deviation respectively of an agent's prior density.

4. Experimental results

This section discusses the experimental results from applying the proposed distributed decision-making model on the simulated environment for different scenarios consisting of the three events: *noEvent*, *fire* and *firework*. In particular, we are interested in understanding how the light-agent and heat-agent infer the environmental states, independently and after mutual communication.

Figure 3 shows the inference by each agent independently for three observation points: $x = 0$, $x = 0.25$ and $x = 1.25$, representing *noEvent*, *firework* and *fire* respectively. The plots show how each agent's belief converges to a particular value of x. *Time* in these plots refers to the duration of time an agent requires to analyze its observation and for the responses (activities) to settle down. The agent finds the most likely value of x by minimizing the free energy. Two prediction errors are involved in the simulation: ε_φ is the difference between observation and its expectation if $x = \mu$, and ε_p is the difference between the belief and the prior expectation.

There is a conflict in the event of *firework* when the light-agent believes it to be a *fire*. The light-agent, however, does not realize its inferred cause is incorrect as there is no prediction error because the light intensity due to *fire* and *firework* are very similar in its generative model (i.e. they share some values of x), as shown in Figure 4.

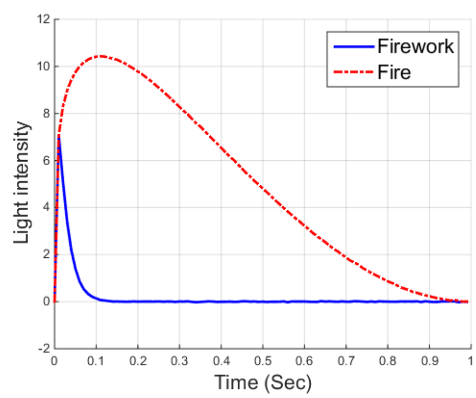

Figure 4: Observations of light-agent for *firework* (x=0.25) and *fire* (x=4) are shown. The initial duration of these events generate the same observation and the agent fails to distinguish between them.

Having the model of heat-agent, the light-agent anticipates the sampling frequency of heat-agent to increase to the range that it should be for the case of *fire* ($x = 4.1$). However, it is surprised as the heat-agent's behavior does not match its expectation. Light-agent initiates communication to minimize its prediction error. The results are shown in Figure 5. It can be seen that the light-agent revised its belief for *firework* since the heat-agent is more confident about its inference (based on the precision, σ_{p2}). Since communication occurs both ways, the belief of heat-agent is slightly increased. However, it still remains in the range of *firework*. Communica-

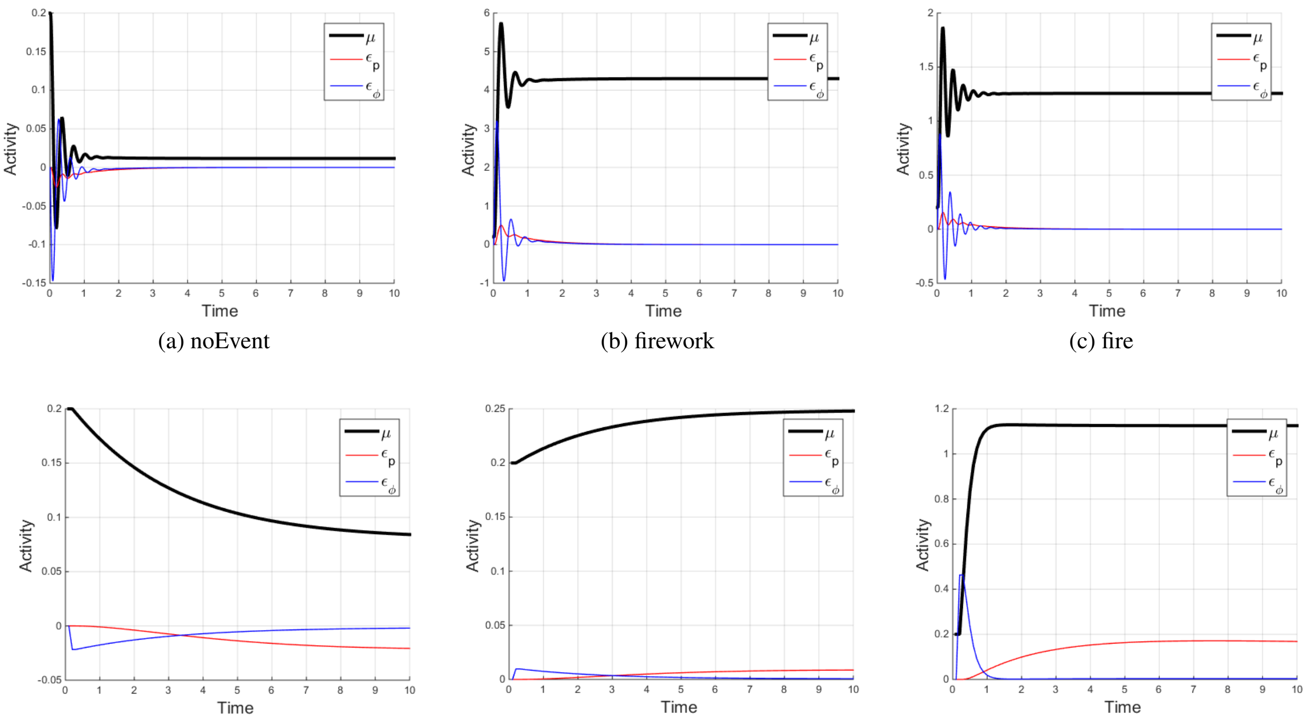

(a) noEvent (b) firework (c) fire

(d) noEvent (e) firework (f) fire

Figure 3: Inference of two agents independently, for a sample of each situation. (Top) Light-agent's inference. (Bottom) Heat-agent's inference. For $firework$, the light-agent (b) converged to $\mu = 4.1$ which is in the range of $fire$. That is, the light-agent made an error in predicting $firework$.

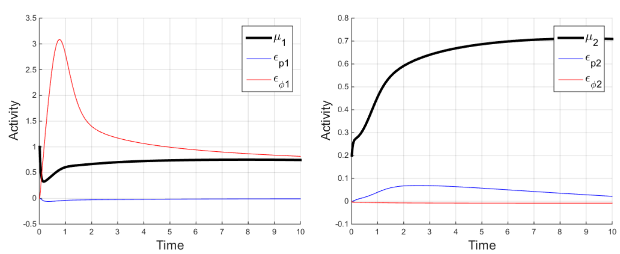

Figure 5: Inference of light-agent (left) and heat-agent (right) regarding $firework$, after communication. The light-agent's inference is improved (it is in the range of $firework$).

tion occurs both ways because the conflict is in the minds of both agents (i.e. the heat-agent also did not predict the message from light-agent and is surprised). The agents continue exchanging messages until the conflict is resolved.

With time, the light intensity due to $fire$ and $firework$ start to differ. Temperature due to heat changes slower than light. Based on these, we construct a scenario of the events $\{noEvent, firework, noEvent, fire\}$, each for 100 seconds duration. Light intensity and temperature observations are shown in Figure 6. The final inferences (after settling down), independently and after mutual communication, are shown in Figure 7. Before communication, the agents fail in two ways: 1) when $firework$ starts, the light-agent infers the cause of its observation incorrectly as $fire$ ($\mu_1 > 1$), and 2) the heat-agent infers the cause of its observation as $fire$ with a significant delay (at time 340, when the $fire$ had started at 300). Both the issues are resolved after communication and their predictions are in the correct ranges. The incorrect inference by the light-agent is resolved as discussed in the current section just after Figure 4. The delay for heat-agent is resolved as follows. The light-agent detects the change earlier and increases its sampling frequency. The heat-agent is surprised by this unexpected change in light-agent's behavior as the former has not detected any significant change in temperature yet. So the heat-agent initiates communication asking the light-agent for the cause of its change in behavior (i.e. the heat-agent samples the light-agent's internal model to minimize its prediction error). The light-agent responds by informing about a significant change in its belief. The conflict is resolved via communication since the heat-agent has learned to trust the light-agent in this situation where the light-agent has high precision (low variance).

5. Conclusions

A novel computational model of distributed decision making is proposed. We show that communication helps a community of predictive coding agents, each limited in its sensori-

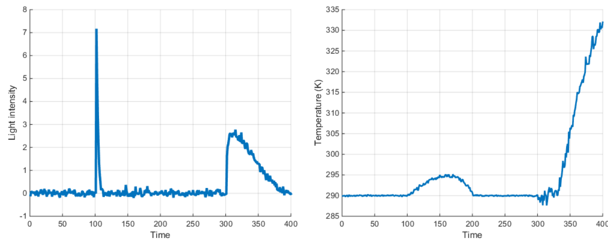

Figure 6: Light intensity (left) and temperature (right) for the simulated scenario of $\{noEvent, firework, noEvent, fire\}$.

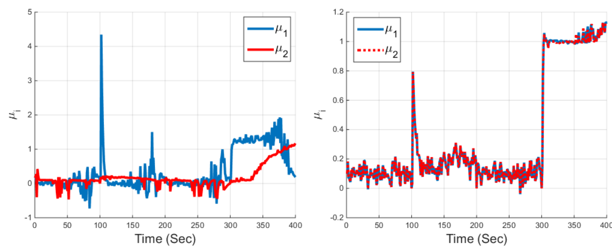

Figure 7: Inference without (left) and with (right) communication. The blue and red lines show the belief of light and heat agents respectively. Conflicts are resolved after communication. [Best viewed in color.]

motor system, to come up with a decision quickly and accurately regarding the state of their shared environment which is not possible for any agent operating independently. The key to this efficiency and accuracy is communication which initiates when a conflict is detected in the mind of an agent due to an error in predicting the other agent's behavior. The proposed model can be scaled to a large number of predictive coding agents operating in a shared environment.

References

Adler, M. R., Davis, A. B., Weihmayer, R., & Worrest, R. W. (1998). Conflict-resolution strategies for nonhierarchical distributed agents. In *Distributed Artificial Intelligence* (pp. 139–161). Elsevier.

Banerjee, B., & Dutta, J. K. (2013, Dec.). A predictive coding framework for learning to predict changes in streaming data. In *IEEE International Conference on Data Mining Workshops* (pp. 497–504). Dallas, TX.

Banerjee, B., & Dutta, J. K. (2014). SELP: A general-purpose framework for learning the norms from saliencies in spatiotemporal data. *Neurocomputing: Special Issue on Brain Inspired Models of Cognitive Memory*, *138*, 41–60.

Bogacz, R. (2015). A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*.

Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*.

Chandler, D., & Munday, R. (2011). *A Dictionary of Media and Communication*. OUP Oxford.

Chen, Y., Murata, S., Arie, H., Ogata, T., Tani, J., & Sugano, S. (2016). Emergence of interactive behaviors between two robots by prediction error minimization mechanism. In *Development and Learning and Epigenetic Robotics, 2016 Joint IEEE International Conference on* (pp. 302–307).

Deneve, S., & Pouget, A. (2004). Bayesian multisensory integration and cross-modal spatial links. *Journal of Physiology-Paris*, *98*(1-3), 249–258.

Di Paolo, E., & De Jaegher, H. (2012). The interactive brain hypothesis. *Frontiers in Human Neuroscience*, *6*.

Foerster, J. N., Assael, Y. M., de Freitas, N., & Whiteson, S. (2016). Learning to communicate to solve riddles with deep distributed recurrent Q-networks. *arXiv preprint arXiv:1602.02672*.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *360*(1456), 815–836.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, *11*(2), 127–138.

Friston, K., Daunizeau, J., & Kiebel, S. J. (2009). Reinforcement learning or active inference? *PloS One*, *4*(7), e6421.

Friston, K., & Frith, C. (2015b). A duet for one. *Consciousness and Cognition*, *36*, 390–405.

Friston, K., & Frith, C. D. (2015a). Active inference, communication and hermeneutics. *Cortex*, *68*, 129–143.

Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2013). The anatomy of choice: active inference and agency. *Frontiers in Human Neuroscience*, *7*, 598.

Müller, H. J., & Dieng, R. (2000). *Computational Conflicts: Conflict Modeling for Distributed Intelligent Systems*. Springer Science & Business Media.

Oleary, D. E. (1999). Identifying conflicting probabilistic knowledge in multiple knowledge base systems. In *Agents Conflicts-Papers from the AAAI Workshop-Technical Report ws-99-08* (pp. 7–13).

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neurosci.*, *2*, 79-87.

Russell, S., & Norvig, P. (2002). *Artificial Intelligence: A Modern Approach* (2nd ed.). Prentice Hall.

Schneeweiss, C. (2012). *Distributed Decision Making*. Springer Science & Business Media.

Sobieska-Karpińska, J., & Hernes, M. (2014). Identification of the knowledge conflicts' sources in the architecture of cognitive agents supporting decision-making process. In *Computer Science and Information Systems (FedCSIS), 2014 Federated Conference on* (pp. 1157–1162).

Sukhbaatar, S., Fergus, R., et al. (2016). Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems* (pp. 2244–2252).