

# The intentional stance as structure learning: a computational perspective on mindreading

Haris Dindo<sup>1</sup> · Francesco Donnarumma<sup>2</sup> · Fabian Chersi<sup>3</sup> · Giovanni Pezzulo<sup>2</sup>

Received: 16 December 2014 / Accepted: 24 June 2015 / Published online: 14 July 2015  
© Springer-Verlag Berlin Heidelberg 2015

**Abstract** Recent theories of mindreading explain the recognition of action, intention, and belief of other agents in terms of generative architectures that model the causal relations between observables (e.g., observed movements) and their hidden causes (e.g., action goals and beliefs). Two kinds of probabilistic generative schemes have been proposed in cognitive science and robotics that link to a “theory theory” and “simulation theory” of mindreading, respectively. The former compares perceived actions to optimal plans derived from rationality principles and conceptual theories of others’ minds. The latter reuses one’s own internal (inverse and forward) models for action execution to perform a look-ahead mental simulation of perceived actions. Both theories, how-

ever, leave one question unanswered: how are the generative models – including task structure and parameters – learned in the first place? We start from Dennett’s “intentional stance” proposal and characterize it within generative theories of action and intention recognition. We propose that humans use an intentional stance as a learning bias that sidesteps the (hard) structure learning problem and bootstraps the acquisition of generative models for others’ actions. The intentional stance corresponds to a candidate structure in the generative scheme, which encodes a simplified belief-desire folk psychology and a hierarchical intention-to-action organization of behavior. This simple structure can be used as a proxy for the “true” generative structure of others’ actions and intentions and is *continuously* grown and refined – via state and parameter learning – during interactions. In turn – as our computational simulations show – this can help solve mindreading problems and bootstrap the acquisition of useful causal models of both one’s own and others’ goal-directed actions.

---

The research leading to these results has received funding from the Human Frontier Science Program, Grant RGY0088/2014 to GP, and the European Union Seventh Framework Programme (FP7/2007-2013) under Grant Agreements FP7-270108 (Goal-Leaders) to GP and 604102 (Human Brain Project) to FC.

---

✉ Giovanni Pezzulo  
giovanni.pezzulo@istc.cnr.it

Haris Dindo  
haris.dindo@unipa.it

Francesco Donnarumma  
francesco.donnarumma@istc.cnr.it

Fabian Chersi  
f.chersi@ucl.ac.uk

- <sup>1</sup> RoboticsLab, Polytechnic School (DICGIM), University of Palermo, Viale delle Scienze, Ed. 6, 90128 Palermo, Italy
- <sup>2</sup> Institute of Cognitive Sciences and Technologies, National Research Council, Via San Martino della Battaglia 44, 00185 Rome, Italy
- <sup>3</sup> Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, UK

**Keywords** Intentional stance · Generative model · Online learning · Structure learning · Mindreading

## 1 Introduction

In our daily activities, for instance when we drive or walk on a crowded street, we systematically and automatically predict what others will do next. More complex forms of social interaction and joint actions require understanding the intentions, beliefs, and preferences of other people, not just predict their next movement, and we excel in all these activities. It is often argued that a hallmark of human cognition is the ability to build theories of others’ minds and to use them to realize collaborative (or competitive) activities with shared goals and

intentions [87]. Still, it is unclear how we solve complex computational problems arising in our everyday social scenarios in real time, such as movement prediction, gaze following, action understanding, intention recognition, belief recognition and tracking, planning what to do next, forming shared representations, and planning how to achieve long-term goals in cooperation or competition with our conspecifics. Clearly, this list includes problems that have different levels of complexity, pose different challenges to social actors, and have been linked to different brain mechanisms [36, 38].

In the rest of this paper we offer a computational account of how such abilities can be bootstrapped. In Sect. 2, we offer a unified perspective on various problems of social interaction using a *generative approach* in which social actors build models (of different complexity) of themselves and their co-actors and use them as an “interaction engine” [60]: a universal resource to understand and predict the actions of others and to plan social interactions. To support our analysis, we discuss examples of how generative models can support mindreading inferences in two different traditions (teleological vs. action-based). In Sect. 3, we ask how generative models used in social interaction are learned in the first place. We propose that the learning process is facilitated by an intentional stance (i.e., the assumption that co-actors are intentional agents), which can be interpreted in terms of fixing a *metastructure* of the proposed generative model and learning its components as the data are collected. In Sect. 4, we present a hierarchical and adaptive probabilistic general model that embeds the idea of intentional stance to facilitate structure and parameter learning.

## 2 A generative approach to mindreading

In principle, we could predict and understand the goals of others using associative mechanisms, which capture regularities between sequences of perceived movements and underlying goals. However, in all but the most trivial social scenarios, the actions of others are highly uncertain, have long-term dependencies, and depend on hidden causes that cannot be directly observed (e.g., the actor’s intentions); all these difficulties are hard to tackle using pure associative mechanisms only [50].

The idea that mindreading problems could be addressed using generative models and probabilistic inference has recently been proposed by various researchers. In one perspective, action and intention recognition are performed using inverse (Bayesian) inference based on a (“teleological”) assumption: that the performing agent is implementing (optimally) a goal-directed action plan (which the Bayesian tries to estimate or recognize) [4, 5]. In another, more “embodied” perspective, the agent reuses the same internal models across action performance and perception; here,

essentially, it is the action performance model (say, an internal model that usually supports a precision grasping action) that is “reused” or “inverted” to support the recognition of a performed action (say, a precision grasping action executed by somebody else) [30, 34, 57, 72, 74]. Although different, both approaches are based on probabilistic generative models, so named because they can both generate and recognize sensory data.

Generative models were first studied in perceptual processing, where they permit recognizing objects or faces by first learning to generate (i.e., imagine or “hallucinate”) them. To do so, the system learns the causal relations between perceptual evidence (say, seeing a cylindrical shape on a table) and its underlying causes (say, the presence of a glass on the table) that cannot be directly observed but must be inferred. A growing body of literature provides efficient computational algorithms for performing inference in probabilistic generative models ranging from state estimation and prediction to learning parameters either from complete or incomplete data [9, 64].

### 2.1 Two views of generative models for action and intention recognition: teleological and action-based

The generative approach has been adopted in many studies in cognitive science and robotics to explain a wide range of mindreading abilities, from at least two different perspectives: *teleological* and *action-based*. The teleological approach [20, 43] proposes that action understanding is based on an inferential process that determines which goals and actions are optimal under certain environmental constraints (or better, the agent’s domain knowledge). Baker and colleagues propose an implementation of the teleological approach that consists in computing an optimal plan  $\pi$  for achieving a goal  $G$  and then solving an *inverse planning* task  $P(G|\pi)$  [5] (but see also [78, 79]). This approach has been extended to explain a variety of phenomena under the umbrella of a “Bayesian theory of mind” [6]; for example, to model belief and desires, a bottom-up inference is made given observed actions [88].

An alternative *action-based* approach stems from motor cognition theories [53] and assumes that internal (inverse and forward) models of one’s own motor system are reused for action estimation or prediction [25, 26, 30, 34, 47, 90, 91]. This view also explains the widespread evidence of motor activation during action perception and prediction and the systematic relations (in terms of development, deficits, and shared neuronal substrate) between our abilities to predict, understand, generate, plan, imagine, and communicate actions [52, 53, 69, 70, 73, 76].

These two approaches link to rival psychological theories (*theory theory* and *simulation theory* of mindreading, respectively) and stress the importance of different kinds of representations and cognitive processes (perceptual rep-

representations and means–ends inference vs. sensorimotor representations and *action simulation*). They have complementary strengths. The former approach focuses principally on high-level mentalistic constructs such as beliefs and desires, eschewing most details of action performance; the latter approach is more detailed at the action level, including, for example, kinematic and dynamic details, but rarely extends to high-level mentalistic constructs. The former approach is (at least in principle) applicable to many cases in which internal (inverse and forward) models implied in sensorimotor control are not available to the perceiving agent [19]. However, it suffers from the difficulties of deriving optimal (or rational) solutions for both its computational burden and the difficulty of obtaining adequate models of (another’s) behavior and costs.

The latter approach overcomes this problem, first, because optimal solutions are not computed, and second, because a model of oneself is used as a proxy for another’s behavior (thereby producing biases in the recognition model) and does not need to be explicitly represented. However, the adequacy of the second approach in explaining complex forms of mindreading such as intention recognition has been questioned because these problems are ill-posed: the same action can be used to satisfy different intentions, which can be indistinguishable based on observation of the proximal action. This problem can be addressed using richer generative models that explicitly model the role of *context* for action disambiguation [30,57]. This would in principle allow for recognizing the proximal goals but also the distal intentions from overt movements. Support for this idea comes from recent studies revealing that the kinematics of proximal actions (say, grasping a glass) are subtly but significantly different depending on the distal intention (say, drinking or throwing the glass) [8]. The neural underpinnings of this model could involve the so-called *action observation network* (ventral premotor cortex, inferior parietal lobule, and superior temporal sulcus) or a wider brain network that also includes a *ventral pathway* [56,57].

There is still controversy on how empirically adequate action-based and teleological approaches are, or whether, instead, they are complementary [37,56]. Here we do not enter into this dispute (but see Sect. 6 for a discussion) but rather focus on the fact that, despite their differences, both theories can be described at the functional level within the same generative approach.

### 3 Structure learning in mindreading

We have discussed how generative architectures can in principle support multiple mindreading abilities; for example, the inference of which intentions, beliefs, and desires, for example, could have produced certain observed movements.

However, the accuracy of the inference depends on the quality of the generative models employed. Thus, one can ask how generative models (encoding “abstract rational theories” in the teleological approach or “sensorimotor experience” in the action-based approach) are learned in the first place.

Acquiring good generative architectures requires solving several problems of increasing difficulty. The simplest case (and obviously the most frequent one) is when both the structure and all the variables of interest are set before learning takes place. The goal is to tune the free parameters of the model (usually conditional probability tables entries or parameters of probability densities) from all the available data. Computational and robotic studies in the action-based tradition use multiple techniques such as supervised or reinforcement learning, active exploration, or imitation learning [13,15,23,24,80]. All these forms of learning pose severe computational challenges and solving them efficiently remains an open problem, although efficient solutions exist for several existing structures [64].

On the other hand, a much more severe problem is that of learning the *structure* of the model to be used, or – in other terms – assessing the existence of causal dependencies between certain variables (say, that the choice of actions depends on intentions) [2,12,85]. Although there is a long history of studies of causality learning in cognitive tasks [45], it is only recently that this kind of learning process has attracted attention in computational studies of decision making [2], motor control [11], and reinforcement learning [44].

In between, there is the possibility of fixing the template of the structure to be used (e.g., a two-level hierarchy of actions or a flat Bayesian network) and to let an algorithm discover the number of variables used (e.g., how many actions?) and their interconnections from data, an approach adopted in this paper. In addition, learning (of the structure, parameters, or both) can be batch (i.e., collect all the samples and run the learning algorithm once) or incremental (i.e., adapt the model for each new data sample), and online (done while the model is functioning in the world) or offline. Needless to say, in most practical applications, the structure of the model is fixed and the learning of its parameters is performed in batch and offline.

As noted by [85], the learning of abstract problem features (such as their underlying structure) is typically framed in a bottom-up way: observable regularities are learned first, and, based on this, increasingly more abstract features are learned using a process of induction or generalization. However, it could be the case that people learn both task parameters and structure at the same time; for instance, they can learn that *action 1* (say, grasping a glass) is often followed by *action 2* (say, bringing the glass to the mouth), and at the same time that this choice of actions is often motivated by an underlying intention (say, to drink). This double problem can be concep-

tualized (in the abstract) as the simultaneous choice of a set of parameters among all their possible combinations and a set of structures among all the possible ones. Note that the two learning processes influence each other such that the choice is between any possible set of parameters *given* a specific structure (repeated for each possible choice of parameters and structures).

This double problem exists also in mindreading contexts, where humans must infer the causal structure underlying another's behavior and simultaneously (learn to) recognize another's actions and intentions. How do humans solve this difficult problem? It has been proposed from a theoretical perspective that humans tend to assume an *intentional stance* toward their conspecifics [27]. In other words, we have a disposition to treat our conspecifics (but also some animals) as rational agents whose behavior is guided by a belief–desire causal structure rather than, say, to just obey the laws of physics. It is worth noting that the *intentional stance* is just a “stance” from the observer's perspective and does not require assuming the existence of beliefs and intentions in the performing agent's mind; it is sufficient that the performing agent behaves *as if* it had beliefs and intentions.

The *intentional stance* hypothesis is popular in philosophy and social neuroscience [27,40], but it has not been linked to generative models or structure learning problems yet. Here we argue that the *intentional stance* can be considered a bias for structure learning because it provides a good – although not necessarily correct – proxy for the generative structure that the mammal brain uses to generate action. In what follows we discuss from a computational perspective the advantages of using the intentional stance.

### 3.1 Using an intentional stance for learning a generative model

As discussed earlier, acquiring a generative model generally involves two distinct learning approaches: we can either fix the structure and learn model parameters from the data using classical machine learning algorithms, or we can try to simultaneously learn both the structure and the parameters from the data. In the former approach, we posit the structure of the model, given our intuitive understanding of the problem, and we infer its parameters such as to maximize a global utility function (for instance, the classical maximum likelihood algorithm provides a point estimate of the parameters that maximize the likelihood of the data for that given model). Structure learning, on the other hand, makes few or no assumptions regarding the structure of the model that should be inferred from the data alongside its parameters. Intuitively, in the absence of other information, any algorithm should traverse the space of *all possible models* and select the one that provides the best explanation of the data according to a fixed metric [59]. However, in all but the most

trivial real-world scenarios, the space of models grows exponentially with the data, which makes structure learning an NP-hard problem from the computational point of view. This is also the case in solving mindreading problems, where in principle observed movements could be caused by any kind of hidden structure.

In practice, however, humans (including children) are extremely accurate in understanding and predicting another's behavior [43]. It has been proposed that the *intentional stance* provides a solution to this problem [27]. From our computational perspective, the intentional stance provides a template of the structure generating the observed behavior. In the language of probabilistic generative models, this structure specifies a series of (in)dependencies between “latent” (i.e., hidden) variables – such as beliefs, proximal goals, distal intentions, and abstract actions – and “observed” ones – such as the perceived behavior of others. In other words, the structure specifies the causal relations between hidden and observable variables and their evolution in time.

It is worth noting that computational modelers have implicitly incorporated an intentional stance (at different levels in teleological and action-based approaches) in their architectures for mindreading. For example, the theory of mind models proposed by [5,6] incorporate knowledge of folk-psychology concepts and specifically the assumption that actions depend on beliefs and desires. Models inspired by action simulation theories [26,30,57] implicitly assume a hierarchically organized intention-to-action structure [48,66,67], pointing to the fact that action sequences achieve distal goals; that low-level actions depend on high-level actions; and that, depending on the context, the same action can have different end effects and achieve different goals.

As intuitive as they might seem, these assumptions play a key role in the models because they translate into (in)dependencies between variables that constrain and greatly simplify inference and learning. These models thus exemplify the advantages of using the intentional stance for determining the structure of mindreading problems. Note that, although different computational proposals incorporate different nodes and concepts (e.g., not all models incorporate utilities or motor primitives) and address different aspects of mindreading, all make use of some form of prior structural information incorporating the fact that mentalistic concepts and intention-to-action hierarchies guide behavior.

### 3.2 Using an intentional stance to bootstrap the learning of task parameters

All the aforementioned models assume a fixed structure and task parameters and principally tackle problems of (action, intention, or belief) inference. None of them directly addresses the issue of how structural information conveyed by an intentional stance supports learning task parameters

(e.g., the number and identity of actions and intentions). Suppose, for example, that a child observes an adult performing an unknown action (e.g., assembling some furniture using a screwdriver). By assuming an intentional stance, the child postulates that the adult's actions are motivated by some intention; however, both the actions and intentions are novel to the child. To successfully understand the adult's behavior, the child must acquire new action and intention information into her generative model. From our computational perspective, this can be considered a learning of new parts of the generative models and task parameters (e.g., number and identity of action and intentions, along with associated parameters that describe their dynamics). We argue that an intentional stance can also help in our example, when observed actions and intentions are novel for the perceiving agent. The initial structure provided by the intentional stance makes it possible to recognize the means–ends structure of the action, assess which parts (models and parameters) are missing, and bootstrap the acquisition of new parts and parameters of the generative model.

#### 4 An adaptive generative model for mindreading

The starting point of our model is the assumption that rational agents' actions are governed by (hidden) intentions. In other words, intentions help explain the behavior of our conspecifics as perceived by an observer. This entails a hierarchical structure, wherein observable action features (e.g., angles of human joints or their derivatives) reside at the lowest level, and actions – intended as a temporal sequence of action features – reside at a higher level (e.g., reaching for and grasping a glass can be seen as a sequence in low-level feature space). What we postulate is that an intention of the agent can be represented as a sequence of actions at this highest level, leading to a distal goal (e.g., reaching for and grasping a glass for drinking). Thus, we do constrain the overall structure of our model based on hints from folk psychology. What we do not constrain is the overall number of states at each level in the model or the connections between said states. In other words, we let the model grow as it observes new data: new states are added at both levels. If we had to leave the structure totally unconstrained (as in classical structure learning – see [85]) – in addition to estimating the states and their interconnection, then we should also tackle the problem of figuring out the “best” structure from the data, a problem known to be NP-hard (see [16]).

Thus, our contribution can be seen from two different points of view: conceptual and computational. At the conceptual level we acknowledge that “intentional stance” makes it possible to derive an explanation from a series of events (caused by another agent). At the computational level we translate this insight into a tractable stochastic model. How-

ever, we would like to stress that the model does not explain the whole spectrum of agent behavior (which would require modeling, for example, beliefs, desires, and emotions). Here we focus on a restricted – yet still significant – example, namely, that of causes of actions at two different temporal scales: proximal vs. distal intentions, postulating a hierarchical two-level structure that grows with data. Thus, any agent behavior can be explained in terms of its proximal/distal effects. This simplifies learning of the model as well as the problem of inference. In a totally unconstrained model, one should postulate the number of levels, the number of states for each level, and their interconnections. One approach could be that of testing all the plausible combinations via a Bayesian model comparison mechanism (intractable for most scenarios except simple ones). Here we adopt the intentional stance heuristics in which the structure is fixed. Such a structure might not be “optimal” in terms of Bayesian structure learning, but nevertheless it produces an acceptable fit.

In what follows we discuss in detail how this process works in practice within an adaptive generative model, called a *growing hierarchical dynamic Bayesian network* (GHDBN), which makes it possible to infer a model's structure and parameters from direct experience *continuously* and *incrementally*.

##### 4.1 GHDBN


The GHDBN is an adaptive (i.e., growing) model that adapts its internal organization when it observes a new behavior while retaining parts of its general structure. We use the adaptive process of the GHDBN as a metaphor of the child's learning process in the previous example. Here the intentional stance fixes the initial *metastructure* of the model in terms of hidden and observed variables and their mutual interdependencies. However, the cardinality (e.g., the number of distinct states) of each variable is not fixed in advance; rather, the model *learns* the best parameterization according to its own experience in the world and does it continually, as new data are collected.

In what follows we provide an illustrative example of a simplified mindreading scenario, roughly equivalent to those usually modeled by other action-based models, in which nodes correspond to observations, simple and complex actions, and distal intentions (Fig. 1). The intentional stance here corresponds mainly to the assumption that any observable behavior of an entity is governed by latent variables representing that entity's hidden proximal goals and distal intentions.

##### 4.2 GHDBN model

From a computational point of view, a GHDBN is an extension of the well-known probabilistic formalism of

**Fig. 1** Growing hierarchical dynamic Bayesian network. The *model* represents hierarchical stochastic processes in which the number of hidden states and their connections in unknown and must be estimated from the observed data as the model acts in the world – hence the attribute “growing”



dynamic Bayesian networks (DBNs) used to model temporal sequences of data that depend on hidden variables. DBNs are characterized by their structure and parameters, where the former specifies the number of variables of the networks, their domain, and conditional dependencies, while the latter quantify their conditional probability distributions.

More formally, a DBN is defined by a set of  $N$  random variables  $\mathbf{X} = \{X^{(1)}, X^{(2)}, \dots, X^{(N)}\}$  and a pair  $\{BN^P, BN^I\}$ , where  $BN^P$  represents the prior distribution over  $X$  at time 1,  $P(\mathbf{X}_1)$ , and  $BN^I$  is a two-slice temporal Bayesian network that defines  $P(\mathbf{X}_t | \mathbf{X}_{t-1})$  by means of a directed acyclic graph (DAG) as follows:

$$P(\mathbf{X}_t | \mathbf{X}_{t-1}) = \prod_{i=1}^N P(X_t^i | Pa(X_t^i)), \quad (1)$$

where  $X_t^i$  is the  $i$ th node at time  $t$  and  $Pa(X_t^i)$  are the parents of  $X_t^i$  in the graph. The parents of a node,  $Pa(X_t^i)$ , can either be in the same time slice or in the previous time slice, i.e., we assume the model is first-order Markov [63].

In general, some variables of the model can be directly observed, while others are hidden and thus can only be inferred from the evolution of observed quantities. Let us denote by  $\mathcal{X}_t$  the set of hidden variables at time  $t$  and by  $\mathcal{Z}_t$  the set of observed variables at the same time step. Any probabilistic inference performed in DBNs, including state estimation, prediction, learning, and model selection, entails computing the posterior distribution over hidden states, given the sequence of observations (a quantity also known as *belief* in Bayesian inference) [63].

Our model, GHDBN, is a hierarchical DBN that postulates a two-level stochastic process, where the higher level induces a sequence of configurations at the lower level, which can be directly observed via a noisy observation process. Observable action features (e.g., angles of human joints or their derivatives) reside at the lowest level, and complete actions – intended as a temporal sequence of action features – reside at a higher level (e.g., reaching for a glass can be seen as a sequence in the low-level feature space). What we postulate is that an agent’s intention can be represented as a sequence of actions at this highest level, leading to a *distal goal* (e.g., reaching for and grasping a glass for drinking; see Fig. 1).

In terms of a probabilistic structure (Fig. 1), the model can be described by two discrete hidden variables,  $X^H$  and  $X^L$ , and two auxiliary binary variables,  $E^H$  and  $E^L$ , representing *absorbing* states at each level of abstraction.<sup>1</sup> Finally, observations are drawn from a multivariate Gaussian distribution represented by the variable  $Y$ .

In an action observation scenario, discrete states of the variable  $X^H$  represent abstract actions (e.g., reaching for, grasping, bringing to mouth), and a path at this level leads to a distal goal (such as drinking). Each abstract action is in turn represented as a (stochastic) path in the space of possible low-level configurations, accounting for the vari-

<sup>1</sup> These are the binary variables used to enforce transition constraints between the two levels in the hierarchy, namely the *blocking constraint* (i.e., the high-level state cannot change before the low-level sequence has terminated, that is,  $E_t^L = 0 \Rightarrow X_t^H = X_{t-1}^H$ ) and the *termination constraint* (i.e., the high-level state cannot terminate before the low-level sequence has terminated, that is,  $E_t^L = 0 \Rightarrow E_t^H = 0$ ).

ability in the action execution and leading to a proximal goal.<sup>2</sup> In other words, discrete states of the low-level variable  $X^L$  represent clusters in a problem-dependent topological space of features, a sequence of which is a synthetic representation of the complex action itself. Each cluster is modeled via a classical multivariate normal distribution, and the whole feature space can be conveniently approximated via a mixture of Gaussians used to draw observations  $Y$ .


While in our model the overall structure is fixed (in terms of hidden and observable placeholder variables and their causal relations), the cardinality of each variable is not known a priori. Note that our approach is different from most structure learning studies in machine learning where the structure of the model itself, together with its parameters, is inferred from the data. As discussed earlier, the general structure learning problem is very hard but can be finessed by assuming initial biases in the otherwise unconstrained learning process [85]. In keeping with our conceptualization of the intentional stance, the model is initialized by postulating a two-level hierarchical structure with unknown states that can “grow” as new data are observed. In this respect, our model bears some similarity to the recent trend of Bayesian nonparametric methods [84], although the statistical machinery involved is radically different.

In the next sections we briefly sketch our computational approach; see [28, 29] for a more detailed description of the algorithms used and experimental results in the fields of imitation learning and human–robot interaction.

## 5 Learning the model

Structure learning includes several problems related to graph construction and to the definition of the number of states and connections between them. Within the GHDBN, the graph is fixed in terms of the number of stochastic variables, their domains, and their interconnections. Our goal is to find an unsupervised method to learn the number of states (i.e., behaviors at both levels of abstraction) and their connections. The problem is challenging since we need to learn the structure of both  $X^H$ 's and  $X^L$ 's states, that is, to discover/modify/remove complex actions (states of  $X^H$ ) and possible transitions between them, and discover/modify/remove low-level events (elementary behaviors, or states of  $X^L$ ) and possible transitions between them.

<sup>2</sup> This idea is consistent with neural evidence for the way humans represent actions, which ranges from simple kinematic acts to more complex assemblies of goal-directed behaviors [48].



**Fig. 2** A 2D topological map created from observations taken from a simple random walk algorithm. *Points* represent observations

### 5.1 Incremental structure learning

Starting with the low-level variable  $X^L$ , we claim that it should reflect the spatial structure of the feature space discretization. In other words, transitions between basic behaviors are only allowed if the corresponding nodes of a topological map over the feature space are connected. In the present work, the observation space is clustered incrementally using the instantaneous topological map (ITM) algorithm [54, 89], which provides a discrete representation of the continuous feature space in the form of a graph where feature space regions are represented as nodes and adjacent regions are connected by edges.<sup>3</sup> The topological map is updated at every new observation  $O_t$  in order to minimize the number of nodes as well as the model distortion. The overall complexity is linear in time and memory with respect to the number of nodes.<sup>4</sup> Figure 2 shows an example of a topological map created by clustering observations taken from a simple random walk algorithm.

Each topological update corresponds to an update of the conditional probability tables (CPTs) in the model. Adding (removing) a node causes an increase (decrease) in the number of  $X^L$  states, corresponding to adding (removing) a column to the prior table  $\pi^{L|H}$ , adding (removing) a column to the end-of-sequence probability table  $P_{end}^{L|H}$ , and adding (removing) a row and a column to each of the transition probability subtables of  $a^{L|H}$ . When an edge is added (removed), a nonzero (zero) probability of the corresponding element

<sup>3</sup> As an example, in a monodimensional feature space representing the velocity, the basic behaviors (nodes) could be as follows: *still* (null value), *walking* (medium value), *running* (high value).

<sup>4</sup> Indeed, a comparison of ITM against well-known self-organizing maps [58] and growing neural gas [39] is performed in [54] and it is shown that ITM outperforms the others even in high-dimensional feature spaces.

in each of the subtables of the  $a^{L|H}$  CPT is set, where a zero value corresponds to an impossible transition between the considered states. To reduce the computational efforts in parameter updating, only the nonzero parameters of the  $a^{L|H}$  CPT are considered. Thus, the number of parameters to be updated in the  $a^{L|H}$  CPT is approximatively  $N \times M \times Mav$ , where  $Mav$  is the average number of neighbors per node. For example, given a topological map with 58 nodes in a three-dimensional feature space with a resolution of  $\epsilon = 0.36$ , the average number of neighbors for each node is 7.9 (see [54]).

In general, it is not always correct to assume that the state transitions on the neighborhood between regions of the feature space can be constrained. In action recognition, however, we are often measuring features (which somehow depend on some physical laws) in a continuous domain. Passing from one state to another one far away in the map should be impossible because changes in observations are usually gradual (if analyzed in a short time window) and correspond to a series of state transitions.

The structure learning of  $X^H$  is more complicated because this variable is not directly observed. The GHDBN could be imagined as many growing hierarchical Markov models (GHMMs) (as the number of  $X^H$  states) unified in a single structure. Having a sequence of observations, it is possible to calculate a likelihood score for each of the submodels to determine the most probable high-level state that has generated the observations. The solution used here is based on measuring the *likelihood ratio test* between the likelihood of the sequence, conditioned on a new  $X^H$  state (representing the current sequence of observations), and the likelihoods conditioned on existing states [49]. If this score is below a (fixed) threshold, we are possibly observing a complex action that has not yet been modeled, so a new  $X^H$  state will be added and its relative parameters estimated. Moreover, the scores for all  $X^H$  states are compared to each other. If the scores of two states differ by a quantity smaller than a preset threshold, then these states are likely to represent a similar complex action and are merged into a unique state.

## 5.2 Parameter learning

Once the structure of a model has been learned, the second problem is how to learn the entries of the various CPTs that will be used as the transition model. The classical approach to parameter learning in probabilistic graphical models is the expectation-maximization (EM) algorithm. Several variants of EM exist, and here we classify an EM algorithm as *batch* or *online*, emphasizing how the whole set of observations is computed to learn the model. EM algorithms basically iterate to maximize a function. Batch EM variants, unfortunately often called online, have been proposed. Two examples are the incremental EM [65] and the stepwise EM [61]. The two algorithms differ in the way they incorporate new infor-

mation, but actually these are not purely online algorithms. However, the main problem is that the more sequences that are observed, the more significant memory and computational efforts are needed. Batch EM algorithms consider the whole set of observations in each iteration, so they also require the initial availability of the whole set of observations (we should collect the data of all the actions we must model before starting the learning phase). Although batch EM ensures convergence at least to a local maximum of the observed data likelihood function, it results in slow learning; moreover, the learned model does not adapt, hence making it inapplicable in real-world applications.

We need a purely online EM algorithm that updates the parameters by processing only the current observation. Several purely online EM algorithms exist, but the majority encompass both the stochastic gradient algorithm and the EM algorithm. An example, used for parameter estimation in Bayesian networks, is the EM( $\eta$ ) algorithm [7], later improved by voting EM [18]. A similar version has been implemented for DBNs (see [17]).

Let  $X_i$  be a node of a generic DBN, and let  $Pa_i$  be the set of the parents of  $X_i$ . An entry in the CPT of the variable  $X_i$  is

$$\theta_{ijk} = P(X_i = x_i^k | Pa_i = Pa_i^j). \quad (2)$$

Thus, the updating rule is as follows:

$$\theta_{ijk}^T = (1 - \eta)\theta_{ijk}^{T-1} + \eta \left[ \frac{P(x_i^k, Pa_i^j | y_t, \theta_{ijk}^{T-1})}{P(Pa_i^j | y_t, \theta_{ijk}^{T-1})} \right]. \quad (3)$$


Such an online updating rule is referred to as stochastic learning, with  $\frac{P(x_i^k, Pa_i^j | y_t, \theta_{ijk}^{T-1})}{P(Pa_i^j | y_t, \theta_{ijk}^{T-1})}$  the instantaneous gradient estimate of the optimization problem constrained by  $\sum_k \theta_{ijk} = 1$ . Online EM algorithms exhibit asymptotic behavior toward the real parameter below a certain variance. This variance is proportional to  $\eta$ , and so the estimate will oscillate around the true parameter.  $\eta$  can be viewed as a *forgetting bias* of the learning algorithm. The system forgets the past at an exponential rate proportional to  $\eta$ . However, the oscillation of the parameter makes it possible to get out of the local maxima when the environment changes and the model needs to adapt its parameters again.

## 5.3 Inference in the model

Since our target is learning complex processes in real time, a recursive filter approach is needed as an inferential algorithm so that received data can be processed sequentially. The technique used here is the well-known particle filter (PF) algorithm [3]. PFs are sequential Monte Carlo algorithms



**Fig. 3** Sampling the GHDBN in a topological order



used to represent probability densities with point masses (or *particles*) in state-space models (see [31] for an exhaustive review of Monte Carlo methods). Note that applying PFs to a GHDBN involves sampling the variables in a topological order for each time slice, as in Fig. 3 (i.e., we first sample the variable labeled 1, then sample variable 2 with the evidence of the sampled value of variable 1, and so on).

### 5.4 Learning algorithm

The overall learning algorithm for the GHDBN is a mix of the previous procedures consisting of two separate phases: *observation* and *learning*; the algorithm is schematically described in the Table 1).

Note that ITM starts with two random nodes;  $X^H$  is initialized to have one state (the first action known is the initial observation sequence), and  $X^L$  states reflect the topological map. We suppose that an observation sequence (in the learning process) is complete (i.e., it must represent a complex event from beginning to end).

Figures 4 and 5 show two examples of the GHDBN evolving in time. In the first case we are considering a simple GHDBN composed of one  $X^H$  state (represented as a rectangle) and three  $X^L$  states (represented as dark circles). Values inside the circles show the probabilities at each time, for example, at time  $t-1$   $P(X_{top}^L|X_1^H) = 0.4$ ,  $P(X_{middle}^L|X_1^H) = 0.5$ , and  $P(X_{bottom}^L|X_1^H) = 0.1$ . In this example, we have

only one high-level state; thus,  $P(X_1^H) = 1$ . Graph arcs (going from left to right) represent possible state transitions. For example, at time  $t-1$ , we could jump from state  $X_{top}^L$  to  $X_{top}^L$  with probability 0.15 and to  $X_{middle}^L$  with probability 0.85; from state  $X_{middle}^L$  to  $X_{top}^L$  with probability 0.10, to  $X_{middle}^L$  with probability 0.40, and to  $X_{bottom}^L$  with probability 0.50. At time  $t$ , a new node is discovered by the ITM clustering algorithm, so a new state is added. This state is supposed to be the neighbor of the node corresponding to  $X_{bottom}^L$ , so a connection is added from  $X_{new}^L$  to  $X_{bottom}^L$  and vice versa (showed by the arcs from time  $t$  to time  $t+1$ ). At time  $t+2$ , ITM decides to delete the node relative to  $X_{top}^L$ , so this state and its connections are removed. Note that transition probabilities change over time owing to the EM learning step.

In Fig. 5, we have a different situation. The observation step in the learning algorithm has detected a new  $X^H$  state (i.e., by comparing its log-likelihood value against a fixed threshold). The diagram now depicts two levels: the background level is the previous  $X_1^H$  state, while the foreground level is the new  $X_2^H$  state. Now high-level state transitions are possible, as represented by the bold arcs. It is worth noting that the low-level PDFs and probability transitions are different from those in the previous example, but the structure is the same.

### 5.5 GHDBN as a model of action understanding

Action understanding is a difficult task as it typically requires predicting intentions hidden from the observation of an action. The GHDBN provides a high level of abstraction and, at the same time, the ability to recognize the basic components of a complex task. To demonstrate the basic principles of the GHDBN model, we have set up a simple experimental scenario consisting of a human manipulating various objects (Fig. 6). Computer vision algorithms continuously segment

**Table 1** GHDBN learning algorithm

*Observe action*



Create a new  $X^H$  state (initializing the CPT as uniform) and learn (with the current sequence of obs.) only the submodel conditioned on it

- a. Compute the likelihoods conditioned on each  $X^H$  state
- b. Compute, for each new state/old state pair, the likelihood ratio
- c. If all these likelihood ratios are greater than a threshold (dependent on the number of parameters and on the length of the sequence), the sequence probably refers to an unknown complex process and the newly added  $X^H$  state should be maintained. Otherwise, remove that state and learn the model normally

*Learn* For each observation

- a. Iterate ITM. Update the  $X^L$  dimension and the structure of the CPTs  $\pi^{L|H}$ ,  $P_{end}^{L|H}$ , and  $a^{L|H}$
- b. If a new  $X^H$  state has been added in pass 1, apply the parameters update rules to all the CPTs conditioned on the new  $X^H$  state, that is, model only the new complex action
- c. Else, apply the parameter update rules to all CPTs

**Fig. 4** An example of updating a GHDBN with one  $X^H$  state. Arcs are directed from left to right in a temporal order



**Fig. 5** A GHDBN with two high-level states. Observations causing the topology changes are the same as in Fig. 4 (notice the same low-level structure, here shown only from  $t - 1$  to  $t + 1$ ) but in this case we suppose that a new  $X^H$  state has been detected

the scene and track the position and velocities of both the hand and all the available objects.

For the task at hand, we have adopted a simple three-dimensional feature vector represented by the following signals: *hand acceleration*, *variation of distance between hand and each object*, and *angle between hand velocity vector and hand-to-object vector*. A distal goal of “moving an object” could be segmented into subactions, for example, *approach object* or *grasp and dislocate object*, and so forth. Each subaction, in turn, is described as a sequence of low-level states and, thus, as a path in the topological map of the feature space. In this case, a simple event could be represented by the node  $[1, -1, 0.01]$  (hand velocity variation, hand–object distance variation, and hand–object angle difference in radians), meaning that the velocity is increasing, the distance between the hand and the object is decreasing, and the angle between the direction of the hand motion and the hand–object direction is small (the hand is moving toward the object). This could be the initial low-level state of the *approaching object* action. The final state could be, for example,  $(-1; -0.1; 0.01)$ , meaning that the velocity is decreasing, as is the distance, but slower than before (the hand is almost near the object).

For instance, suppose we observe an approach–object action. How is learning performed? Initially, the GHDBN has only one default  $X^H$  state, so this action will be assigned to that state (say, with  $id_1$ ). If the system observes a radically different action, for instance *grasp and dislocate object*, a new  $X^H$  state should be created, and this action would have been given  $id_2$ , and so on.

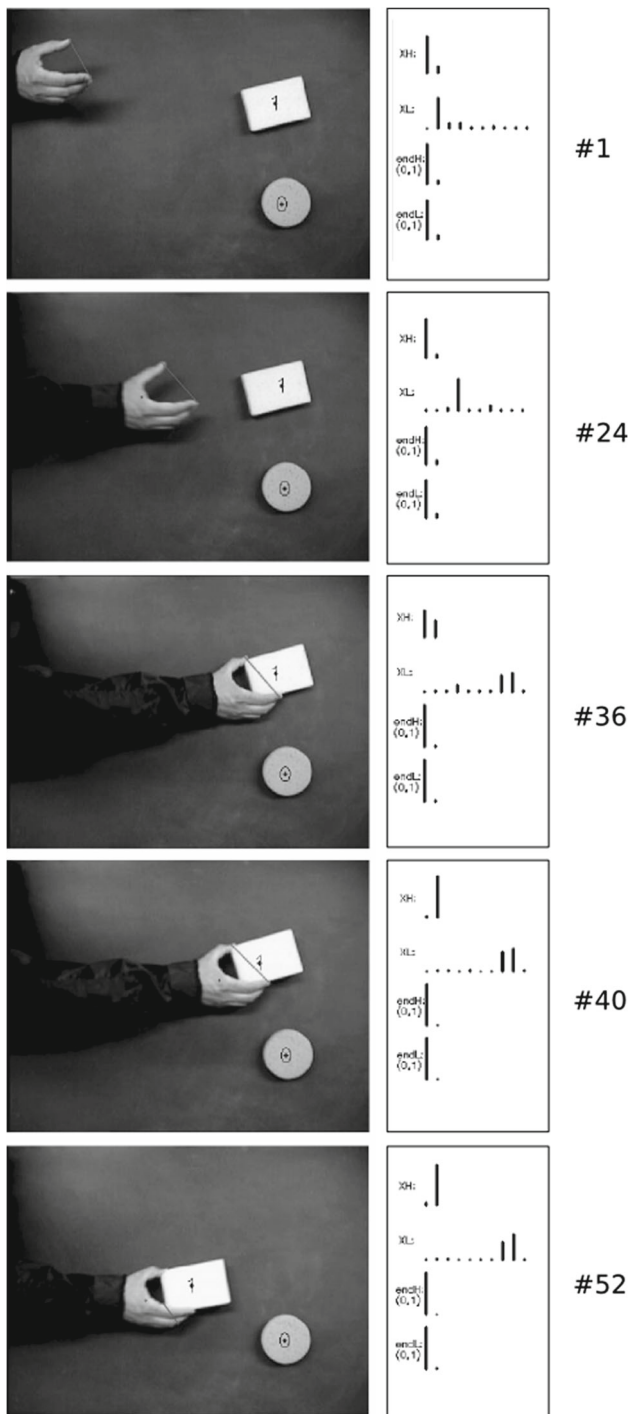
Figure 6 shows an example of action recognition (after the learning of two high-level actions). For each frame, the distributions of  $X^H$ ,  $X^L$ ,  $E^H$ , and  $E^L$  are shown. Let us focus on the evolution of the  $X^H$  state: from frame #1 to frame #24 the most probable state is the first one (corresponding to the approach–object action); at frame #36 the GHDBN observes a possible transition to action 2 (corresponding to the grasp–and–dislocate–object action); after this frame the most probable action is indeed the second one.

Figure 6 shows an example of action recognition (we assume the system had already observed a similar behavior and had learned the model and its parameters).

### 6 Discussion and future works

In this paper we have presented a generative approach to problems of mindreading. We have emphasized that two kinds of generative scheme have been proposed in cognitive science, neuroscience and robotics, which link to “theory theory” and “simulation theory” of mindreading, respectively. While the former compares perceived actions to optimal plans derived from rationality principles and abstract “theories” of mind (i.e., models that describe how they behave), the latter reuses its own internal (forward and inverse) models to perform a look-ahead mental simulation of perceived actions.

Both theories require reliable models to work properly, and their acquisition requires learning both task parameters and structure. Here we argue that an intentional stance helps solve the structure learning process by providing a proxy for the “real” structure used by the brain to derive actions from intentions, and vice versa. We discussed how structure learn-



**Fig. 6** Recognition of a composite action (approach object and grasp and dislocate object). Next to each frame we depict the histogram of hidden GHDBN variables:  $X^H$ , high-level variable (two learned states);  $X^L$ , low-level variable (ten learned states);  $E^H$ , end-of-high-level-sequence binary variable; and  $E^L$ , end-of-high-level-sequence binary variable. From frame #1 to frame #36 the most probable state is the first one (corresponding to the approach-object action); from the frame #40 until the end, the most probable high-level state is the second one (corresponding to the grasp-and-dislocate-object action). Shown also are the activations of low-level states during observation

ing can help the mindreading process, and we provided an example (the GHDBN) of how structure initialization can help in the learning of new parts of the model (e.g., actions and intentions). The example we provided targets the recognition of actions and intentions (equivalent to models of the action-based tradition); it needs to be extended with other components (encoding, for example, probabilistic relations between beliefs, desires, and intentions) to manage more complex cases of mindreading.

Relative to this issue, it is worth noting that both action-based and teleological approaches (implicitly or explicitly) use a concept of intentional stance, but these concepts are slightly different. In the *action-based* view of generative models, a kind of intentional stance is somewhat implicit in the use of oneself as a model of another person, as recognized in the “like-me” hypothesis of [62], which links to both mentalistic processes and emotion attribution. In the *teleological* view of generative models, the intentional stance is an explicit attribution process in which the learning agent can use the *folk-psychology* concepts of desires, beliefs, and intentions (and their relations) as a proxy for the “real” generative structure that the mammalian brain uses to generate action.

As discussed by [27], the success of mentalistic explanations of behavior ensure the survival of folk-psychology concepts, despite the fact that their ontological status is unclear. Ultimately, from a computational perspective, it is only the success of causal inferences that matters, whether or not the underlying model describes real entities. Thus, the successful *use* of an intentional stance does not imply per se that the structure it postulates is correct (or that folk-psychology concepts are adequate to explain cognition and behavior). However, recent neuroscience studies reveal that the idea of a hierarchical generative model is at least a good approximation of how the brain of living organisms solves the intention-to-action and action-to-intention problems [35,48], partially explaining the computational efficacy (although not completely proving the empirical plausibility) of explanations based on the intentional stance.

Although many scholars have focused on whether mentalistic concepts such as “beliefs” and “desires” *exist* in the brain, in the proposed perspective, the only assumption is that the brain encodes dynamic processes at different hierarchical levels and time scales, arranged in a generative scheme in which high-level processes drive low-level ones [55]. In this view, long-lasting brain states (analogous to – although not necessarily identical to – beliefs and intentions) could guide low-level action and behavior dynamics through a generative scheme [33,35]. If this view is correct, the structure postulated by the intentional stance could ultimately boil down to the fact that agentic behavior is better explained in terms of deep architectures having multilevel dynamics, and that certain hidden states are long-lasting and “constitutive” of an

agent (e.g., its beliefs and intentions) and distinguish it from other individuals – which is not generally true of inanimate objects.

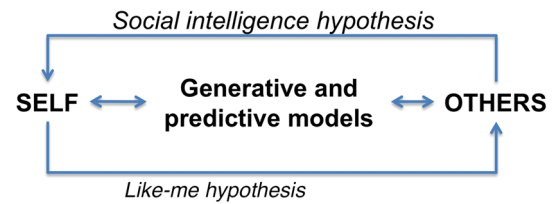
We have proposed that in the same way the brain builds generative models of its sensorium, it could build (hierarchical and generative) models of others' behavior; but rather than building them by a slow inductive process, it could use an intentional stance to avoid costly structure learning. Our computational proposal is an alternative to both the innatist view that a theory of mind cannot be learned from experience and to the associationist view that theory of mind could be acquired using associative rather than generative mechanisms. Furthermore, as it adopts a rapid structural expansion process starting from a fixed template – rather than a slow bottom-up inductive process – it is compatible with evidence that elements of a theory of mind are present early in infancy [20] and guide infants' knowledge acquisition [21].

In what follows, we discuss some additional arguments that stem from our computational approach to mindreading problems and that constitute open objectives for future research.

### 6.1 Using an intentional stance to explain the behavior of living organisms, and beyond

If learning agents tend to assume an intentional stance, it is not implausible that they postulate an “agentive” structure among their very first hypotheses. Indeed, not only is an intentional stance adopted to explain the behavior of humans and other living organisms, but it is often applied to a much larger class of processes (e.g., the movement of nonliving objects such as balls or clouds, or to provide explanations of why computers crash or why it rains).

These considerations suggest that causal reasoning is indeed central to human cognition [46,68]. During infancy, children (like other animals) plausibly try to learn the causal structure of their world. A classical experiment has shown that pigeons can learn spurious causal relations that lead to “superstitious” behavior [83]. During learning, mentalistic explanations for nonliving phenomena are mostly ruled out, probably owing to a combination of poor prediction accuracy and cultural learning. Still, even after significant learning, certain features of humans and other living organisms, such as characteristic traits (e.g., eyelike or smilelike elements), certain kinds of movement (e.g., animal-like), or the predictability of certain patterns of feedback, prompt the causal attribution of folk-psychology structures. This would explain why we attribute intentionality preferentially to objects designed to look like or move as humans or that have complex but predictable structures (such as computers and cars), although we know that the hypothesized structure could not be empirically true (i.e., cars do not have desires, although their designers could).



**Fig. 7** Ontogenesis of social cognition: from individual to social and vice versa

### 6.2 Ontogenesis of social cognition: from individual to social and vice versa

Our computational analysis links to the philosophical problem of explaining the links between individual and social cognition, and the role of attribution processes for passing from the former to the latter, or vice versa. As illustrated in Fig. 7, generative models of self can be used to understand others (similar to the “like-me” hypothesis). Alternatively, generative models (and the intentional stance) could have been initially developed to explain others' behavior and be successively adopted to the self, which would account for a social origin of intelligence. Future research using generative models could help distinguish the two alternatives on empirical bases.

### 6.3 Debate between action-based and teleological approaches

As discussed earlier, both teleological and action-based approaches to mindreading can be cast in terms of generative (Bayesian) models. The former appeals to inverse inference in generative models to explain mindreading in terms of rationality principles [5,20]. The latter uses the idea of generative models to explain the commonalities between the neural circuits of action performance, observation, and prediction; for instance, the action execution model is also used as a recognition model [34,56,57]. Although it is not typically framed in this way, the action-based approach considers issues of optimality in the recognition process. It is often assumed that action planning and execution follow optimality principles and minimize costs due to, for example, effort or time [82,86]. When the same internal models are used for action recognition, the same optimality principles are (implicitly) incorporated into the recognition process.

From a computational viewpoint, these two approaches solve the same set of problems (e.g., action understanding) using partially different sources of information (e.g., motor processes for the action-based approach vs. primarily visual properties of objects for the teleological approach), which could be more or less available in different contexts. This fact suggests that action-based and teleological mechanisms could coexist and cooperate to solve mindreading problems

in a common generative framework, in which the brain essentially uses (and fuses) all the available sources of information. Furthermore, our emphasis on the generative approach does not rule out the possibility of using associative mechanisms in combination with them. Of course, this computational analysis must be substantiated by empirical research, and thus assessing the contribution of these mechanisms in mindreading task is an open challenge for future research.

#### 6.4 Extending the generative approach to joint actions and communication

The computational architecture for mindreading that we described can help in examining how increasingly more complex forms of social interaction, such as those implying joint actions and (linguistic or nonlinguistic) communication, develop on top of the same “interaction engine” [60] for generating and recognizing action.

Indeed, in recent years cognitive psychology and neuroscience have been shifting their attention from the study of passive action observation (and mindreading) scenarios discussed herein to more active forms of interaction. Indeed, mindreading abilities are part of a complex capacity for social interaction (an “interaction engine”) that also includes mechanisms for planning social actions, creating and picking up affordances from others, achieving joint goals, forming shared representations, engaging in social interactions (cooperative or competitive), helping or hindering others, communicating with others, influencing them, and changing their goals.

These processes remain incompletely understood from both empirical and computational perspectives, and their study will likely form the agenda of social psychology and neuroscience for many years to come. If we focus on the somewhat more restricted scenario of “simple” joint actions, and specifically face-to-face interactions such as dancing or moving a table together, consensus is accumulating for the idea that perceptual and action processes (say, recognizing an action and executing the same or a complementary one) largely use shared neural resources [32,41,51,53,77]. As this body of evidence links well to the action-based family of generative models discussed so far, it suggests a possible developmental process and a reuse of computational processes from simpler to more complex social scenarios.

Still, it remains to be understood how production and recognition processes are integrated and orchestrated so smoothly that (despite their computational complexities) joint actions seem to be performed in an effortless manner. A useful starting point for understanding joint action problems in their complexity is the proposal of Sebanz et al. [81] that an architecture for joint action should solve, in an integrated framework, two kinds of problems: high-level (understanding and planning actions to achieve joint

goals) and low-level (prediction and coordination of behavior in real time to achieve such goals). Furthermore, all these processes interact with bottom-up dynamics of interaction, such as the automatic alignment of behavior and mutual emulation [42], which can also influence higher-level processes of goal selection and action prediction [1, 10]. Initial computational proposals have been put forward that explain joint action [74] (see also [14, 22, 72, 75]) and communication [71] using and extending the same generative approach elucidated so far. Of course, elucidating the neurocomputational architecture of mindreading in the context of complex joint actions and communication remains an open objective for future research.

**Acknowledgments** The authors would like to thank Guido Schillaci for useful discussions and background materials.

#### References

1. Aarts H, Gollwitzer P, Hassin R (2004) Goal contagion: perceiving is for pursuing. *J Pers Soc Psychol* 87:23–37
2. Acuna DE, Schrater P (2010) Structure learning in human sequential decision-making. *PLoS Comput Biol* 6(12):e1001003. doi:10.1371/journal.pcbi.1001003
3. Arulampalam MS, Maskell S, Gordon N, Clapp T (2002) A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans Signal Process* 50(2):174–188
4. Baker C, Tenenbaum J, Saxe R (2006) Bayesian models of human action understanding. In: Weiss Y, Schölkopf B, Platt J (eds) *Advances in neural information processing systems 18*. MIT Press, Cambridge, pp 99–106
5. Baker CL, Saxe R, Tenenbaum JB (2009) Action understanding as inverse planning. *Cognition* 113(3):329–349. doi:10.1016/j.cognition.2009.07.005
6. Baker CL, Saxe RR, Tenenbaum JB (2011) Bayesian theory of mind: modeling joint belief-desire attribution. In: *Proceedings of the thirty-second annual conference of the cognitive science society*
7. Bauer E, Koller D, Singer Y (1997) Update rules for parameter estimation in Bayesian networks. In: *Proceedings of the thirteenth conference on uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc, pp 3–13
8. Becchio C, Manera V, Sartori L, Cavallo A, Castiello U (2012) Grasping intentions: from thought experiments to empirical evidence. *Front Hum Neurosci* 6:117. doi:10.3389/fnhum.2012.00117
9. Bishop CM (2006) *Pattern recognition and machine learning*. Springer, New York
10. Blakemore SJ, Frith C (2005) The role of motor contagion in the prediction of action. *Neuropsychologia* 43(2):260–267. doi:10.1016/j.neuropsychologia.2004.11.012
11. Braun DA, Mehring C, Wolpert DM (2010) Structure learning in action. *Behav Brain Res* 206(2):157–165. doi:10.1016/j.bbr.2009.08.031
12. Carroll CD, Kemp C (2013) Hypothesis space checking in intuitive reasoning. In: *Proceedings of the 35th annual conference of the cognitive science society*. Cognitive science society
13. Chella A, Dindo H, Infantino I (2007) Imitation learning and anchoring through conceptual spaces. *Appl Artif Intell* 21(4–5):343–359. doi:10.1080/08839510701252619

14. Chersi F (2011) Neural mechanisms and models underlying joint action. *Exp Brain Res* 211(3–4):643–653. doi:[10.1007/s00221-011-2690-3](https://doi.org/10.1007/s00221-011-2690-3)
15. Chersi F (2012) Learning through imitation: a biological approach to robotics. *IEEE Trans Auton Ment Dev* 4(3):204–214. doi:[10.1109/TAMD.2012.2200250](https://doi.org/10.1109/TAMD.2012.2200250)
16. Chickering DM, Geiger D, Heckerman D et al (1994) Learning Bayesian networks is NP-hard. Tech. rep, Citeseer
17. Cho HC, Fadali SM (2006) Online estimation of dynamic Bayesian network parameter. In: IEEE International joint conference on neural networks, 2006, IJCNN'06, pp 3363–3370
18. Cohen I, Bronstein A, Cozman FG (2001) Adaptive online learning of Bayesian network parameters. University of Sao Paulo, Brasil
19. Csibra G (2007) Action mirroring and action understanding: an alternative account. In: Haggard P, Rosetti Y, Kawato M (eds) *Sensorimotor foundations of higher cognition: attention and performance XXII*. Oxford University Press, Oxford
20. Csibra G, Gergely G (2007) Obsessed with goals: functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychol* 124:60–78
21. Csibra G, Gergely G (2009) Natural pedagogy. *Trends Cogn Sci* 13(4):148–153. doi:[10.1016/j.tics.2009.01.005](https://doi.org/10.1016/j.tics.2009.01.005)
22. Cuijpers RH, van Schie HT, Koppen M, Erlhagen W, Bekkering H (2006) Goals and means in action observation: a computational approach. *Neural Netw* 19(3):311–322. doi:[10.1016/j.neunet.2006.02.004](https://doi.org/10.1016/j.neunet.2006.02.004)
23. Dearden A, Demiris Y (2005) Learning forward models for robotics. In: *Proceedings of IJCAI-2005*. Edinburgh, pp 1440–1445
24. Demiris J, Hayes G (2002) Imitation as a dual-route process featuring predictive and learning components: a biologically plausible computational model. *Imitation in animals and artifacts*. MIT Press, Cambridge, pp 327–361
25. Demiris Y (2007) Prediction of intent in robotics and multi-agent systems. *Cogn Process* 8(3):151–158
26. Demiris Y, Khadhour B (2005) Hierarchical attentive multiple models for execution and recognition (Hammer). *Rob Auton Syst J* 54:361–369
27. Dennett D (1987) *The intentional stance*. MIT Press, Cambridge
28. Dindo H, Schillaci G (2010) An adaptive probabilistic graphical model for representing skills in pbd settings. In: *Proceedings of the 5th ACM/IEEE international conference on human-robot interaction*, pp 89–90
29. Dindo H, Schillaci G (2010) An adaptive probabilistic approach to goal-level imitation learning. In: *Proceedings of the 2010 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp 4452–4457. doi:[10.1109/IROS.2010.5654440](https://doi.org/10.1109/IROS.2010.5654440)
30. Dindo H, Zambuto D, Pezzulo G (2011) Motor simulation via coupled internal models using sequential Monte Carlo. *Proceedings of IJCAI 2011*:2113–2119
31. Doucet A, De Freitas N, Gordon N et al (2001) *Sequential Monte Carlo methods in practice*, vol 1. Springer, New York
32. Fadiga L, Fogassi L, Pavesi G, Rizzolatti G (1995) Motor facilitation during action observation: a magnetic stimulation study. *J Neurophysiol* 73:2608–2611
33. Friston K (2005) A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360(1456):815–836. doi:[10.1098/rstb.2005.1622](https://doi.org/10.1098/rstb.2005.1622)
34. Friston K, Mattout J, Kilner J (2011) Action understanding and active inference. *Biol Cybern* 104(1–2):137–160. doi:[10.1007/s00422-011-0424-z](https://doi.org/10.1007/s00422-011-0424-z)
35. Friston KJ, Daunizeau J, Kilner J, Kiebel SJ (2010) Action and behavior: a free-energy formulation. *Biol Cybern* 102(3):227–260. doi:[10.1007/s00422-010-0364-z](https://doi.org/10.1007/s00422-010-0364-z)
36. Frith CD, Frith U (2006) How we predict what other people are going to do. *Brain Res* 1079(1):36–46. doi:[10.1016/j.brainres.2005.12.126](https://doi.org/10.1016/j.brainres.2005.12.126). <http://www.sciencedirect.com/science/article/B6SYR-4JCCJXS-3/2/ec6adb1fa305f3d6791ac45f451cf63d>
37. Frith CD, Frith U (2008) Implicit and explicit processes in social cognition. *Neuron* 60(3):503–510. doi:[10.1016/j.neuron.2008.10.032](https://doi.org/10.1016/j.neuron.2008.10.032)
38. Frith U, Frith C (2010) The social brain: allowing humans to boldly go where no other species has been. *Philos Trans R Soc Lond B Biol Sci* 365(1537):165–176. doi:[10.1098/rstb.2009.0160](https://doi.org/10.1098/rstb.2009.0160)
39. Fritzsche B et al (1995) A growing neural gas network learns topologies. *Adv Neural Inf Process Syst* 7:625–632
40. Gallagher HL, Jack AI, Roepstorff A, Frith CD (2002) Imaging the intentional stance in a competitive game. *Neuroimage* 16(3 Pt 1):814–821
41. Gallese V, Fadiga L, Fogassi L, Rizzolatti G (1996) Action recognition in the premotor cortex. *Brain* 119:593–609
42. Garrod S, Pickering MJ (2009) Joint action, interactive alignment, and dialog. *Top Cogn Sci* 1(2):292–304. doi:[10.1111/j.1756-8765.2009.01020.x](https://doi.org/10.1111/j.1756-8765.2009.01020.x)
43. Gergely G, Csibra G (2003) Teleological reasoning in infancy: the naive theory of rational action. *Trends Cogn Sci* 7:287–292
44. Gershman SJ, Niv Y (2010) Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol* 20(2):251–256. doi:[10.1016/j.conb.2010.02.008](https://doi.org/10.1016/j.conb.2010.02.008)
45. Gopnik A, Glymour C, Sobel DM, Schulz LE, Kushnir T, Danks D (2004) A theory of causal learning in children: causal maps and Bayes nets. *Psychol Rev* 111(1):3–32. doi:[10.1037/0033-295X.111.1.3](https://doi.org/10.1037/0033-295X.111.1.3)
46. Gopnik A, Meltzoff A (1997) *Words, thoughts and theories*. MIT Press, Cambridge
47. Grush R (2004) The emulation theory of representation: motor control, imagery, and perception. *Behav Brain Sci* 27(3):377–396
48. Hamilton AFdC, Grafton ST (2007) The motor hierarchy: from kinematics to goals and intentions. In: Haggard P, Rossetti Y, Kawato M (eds) *Sensorimotor foundations of higher cognition*. Oxford University Press, Oxford
49. Heckerman D (1998) *A tutorial on learning with Bayesian networks*. Springer, New York
50. Heil L, van Pelt S, Kwisthout J, van Rooij I, Bekkering H (2014) Higher-level processes in the formation and application of associations during action understanding. *Behav Brain Sci* 37(02):202–203
51. Hommel B, Musseler J, Aschersleben G, Prinz W (2001) The theory of event coding (tec): a framework for perception and action planning. *Behav Brain Sci* 24(5):849–878
52. Jeannerod M (2001) Neural simulation of action: a unifying mechanism for motor cognition. *NeuroImage* 14:S103–S109
53. Jeannerod M (2006) *Motor cognition*. Oxford University Press, Oxford
54. Jockusch J, Ritter H (1999) An instantaneous topological mapping model for correlated stimuli. In: *Proceedings of the international joint conference on neural networks*, Washington (US), vol 1, pp 529–534
55. Kiebel SJ, von Kriegstein K, Daunizeau J, Friston KJ (2009) Recognizing sequences of sequences. *PLoS Comput Biol* 5(8):e1000464. doi:[10.1371/journal.pcbi.1000464](https://doi.org/10.1371/journal.pcbi.1000464)
56. Kilner JM (2011) More than one pathway to action understanding. *Trends Cogn Sci* 15(8):352–357. doi:[10.1016/j.tics.2011.06.005](https://doi.org/10.1016/j.tics.2011.06.005)
57. Kilner JM, Friston KJ, Frith CD (2007) Predictive coding: an account of the mirror neuron system. *Cogn Process* 8(3):159–166
58. Kohonen T (1998) *The self-organizing map*. Neurocomputing 21(1):1–6
59. Koller D, Friedman N (2009) *Probabilistic graphical models: principles and techniques*. MIT Press, Cambridge
60. Levinson SC (2006) On the human “interaction engine”. In: Enfield NJ, Levinson SC (eds) *Roots of human sociality: culture, cognition and interaction*. Berg, Oxford, pp 39–69

61. Liang P, Klein D (2009) Online em for unsupervised models. In: Proceedings of human language technologies: the 2009 annual conference of the North American chapter of the association for computational linguistics, pp 611–619. Association for computational linguistics
62. Meltzoff AN (2007) 'like me': a foundation for social cognition. *Dev Sci* 10(1):126–134. doi:[10.1111/j.1467-7687.2007.00574.x](https://doi.org/10.1111/j.1467-7687.2007.00574.x)
63. Murphy KP (2002) Dynamic Bayesian networks: representation, inference and learning. Ph.D. thesis, UC Berkeley, Computer Science Division. <http://www.worldcat.org/oclc/52827959>
64. Murphy KP (2012) Machine learning: a probabilistic perspective. MIT Press, Cambridge
65. Neal RM, Hinton GE (1998) A view of the EM algorithm that justifies incremental, sparse, and other variants. In: Jordan MI (ed) Learning in graphical models. Kluwer Academic Publishers, Norwell, pp 355–368
66. Ognibene D, Wu Y, Lee K, Demiris Y (2013) Hierarchies for embodied action perception. In: Computational and robotic models of the hierarchical organization of behavior. Springer, pp 81–98
67. Pacherie E (2008) The phenomenology of action: a conceptual framework. *Cognition* 107:179–217
68. Pearl J (2000) Causality: models, reasoning, and inference. Cambridge University Press, Cambridge
69. Pezzulo G (2008) Coordinating with the future: the anticipatory nature of representation. *Minds Mach* 18(2):179–225. doi:[10.1007/s11023-008-9095-5](https://doi.org/10.1007/s11023-008-9095-5)
70. Pezzulo G (2011) Grounding procedural and declarative knowledge in sensorimotor anticipation. *Mind Lang* 26(1):78–114
71. Pezzulo G (2012) The interaction engine: a common pragmatic competence across linguistic and non-linguistic interactions. *IEEE Trans Auton Ment Dev* 4(2):105–123
72. Pezzulo G (2013) Studying mirror mechanisms within generative and predictive architectures for joint action. *Cortex* 49:2968–2969
73. Pezzulo G, Candidi M, Dindo H, Barca L (2013) Action simulation in the human brain: twelve questions. *New Ideas Psychol* 31(3):270–290
74. Pezzulo G, Dindo H (2011) What should i do next? Using shared representations to solve interaction problems. *Exp Brain Res* 211(3):613–630
75. Pezzulo G, Donnarumma F, Dindo H (2013) Human sensorimotor communication: a theory of signaling in online social interactions. *PLoS One* 8(11):e79876. doi:[10.1371/journal.pone.0079876](https://doi.org/10.1371/journal.pone.0079876)
76. Pezzulo G, van der Meer MA, Lansink CS, Pennartz CM (2014) Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn Sci* 18(12):647–657
77. Pickering MJ, Garrod S (2007) Do people use language production to make predictions during comprehension? *Trends Cogn Sci* 11(3):105–110
78. Ramirez M, Geffner H (2010) Probabilistic plan recognition using off-the-shelf classical planners. In: Proceedings of AAAI-10. Atlanta, USA
79. Ramirez M, Geffner H (2011) Goal recognition over pomdps: inferring the intention of a pomdp agent. In: IJCAI, pp 2009–2014
80. Schaal S (1999) Is imitation learning the route to humanoid robots? *Trends Cogn Sci* 3:233–242
81. Sebanz N, Bekkering H, Knoblich G (2006) Joint action: bodies and minds moving together. *Trends Cogn Sci* 10(2):70–76. doi:[10.1016/j.tics.2005.12.009](https://doi.org/10.1016/j.tics.2005.12.009)
82. Shadmehr R, Mussa-Ivaldi S (2012) Biological learning and control: how the brain builds representations, predicts events, and makes decisions. MIT Press, Cambridge
83. Skinner BF (1948) Superstition in the pigeon. *J Exp Psychol* 38(2):168–172
84. Teh YW, Jordan MI (2010) Hierarchical Bayesian nonparametric models with applications. *Bayesian Nonparametrics*. Cambridge University Press, pp 158–207
85. Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND (2011) How to grow a mind: statistics, structure, and abstraction. *Science* 331(6022):1279–1285. doi:[10.1126/science.1192788](https://doi.org/10.1126/science.1192788)
86. Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5(11):1226–1235. doi:[10.1038/nn963](https://doi.org/10.1038/nn963)
87. Tomasello M, Carpenter M, Call J, Behne T, Moll H (2005) Understanding and sharing intentions: the origins of cultural cognition. *Behav Brain Sci* 28(5):675–691. doi:[10.1017/S0140525X05000129](https://doi.org/10.1017/S0140525X05000129)
88. Ullman T, Tenenbaum J, Baker C, Macindoe O, Evans O, Goodman N et al. (2009) Help or hinder: Bayesian models of social goal inference. In: Proceedings of neural information processing systems
89. Vasquez D, Fraichard T, Laugier C (2009) Incremental learning of statistical motion patterns with growing hidden markov models. *IEEE Trans Intell Trans Syst* 10(3):403–416
90. Wilson M, Knoblich G (2005) The case for motor involvement in perceiving conspecifics. *Psychol Bull* 131:460–473
91. Wolpert DM, Doya K, Kawato M (2003) A unifying computational framework for motor control and social interaction. *Philos Trans R Soc Lond B Biol Sci* 358(1431):593–602. doi:[10.1098/rstb.2002.1238](https://doi.org/10.1098/rstb.2002.1238)