



# SpeedyIBL: A comprehensive, precise, and fast implementation of instance-based learning theory

Thuy Ngoc Nguyen<sup>1</sup> · Duy Nhat Phan<sup>1</sup> · Cleotilde Gonzalez<sup>1</sup>

Accepted: 22 March 2022 / Published online: 29 June 2022  
© The Psychonomic Society, Inc. 2022

## Abstract

Instance-based learning theory (IBLT) is a comprehensive account of how humans make decisions from experience during dynamic tasks. Since it was first proposed almost two decades ago, multiple computational models have been constructed based on IBLT (i.e., IBL models). These models have been demonstrated to be very successful in explaining and predicting human decisions in multiple decision-making contexts. However, as IBLT has evolved, the initial description of the theory has become less precise, and it is unclear how its demonstration can be expanded to more complex, dynamic, and multi-agent environments. This paper presents an updated version of the current theoretical components of IBLT in a comprehensive and precise form. It also provides an advanced implementation of the full set of theoretical mechanisms, SpeedyIBL, to unlock the capabilities of IBLT to handle a diverse taxonomy of individual and multi-agent decision-making problems. SpeedyIBL addresses a practical computational issue in past implementations of IBL models, the curse of exponential growth, that emerges from memory-based tabular computations. When more observations accumulate over time, there is an exponential growth of the memory of instances that leads directly to an exponential slowdown of the computational time. Thus, SpeedyIBL leverages parallel computation with vectorization to speed up the execution time of IBL models. We evaluate the robustness of SpeedyIBL over an existing implementation of IBLT in decision games of increased complexity. The results not only demonstrate the applicability of IBLT through a wide range of decision-making tasks, but also highlight the improvement of SpeedyIBL over its prior implementation as the complexity of decision features the number of agents increase. The library is open sourced for the use of the broad research community.

**Keywords** Instance-based learning · Cognitive models · Decision from experience · Python instance-based learning library

## Introduction

A cognitive theory is a general postulation of mechanisms and processes that are globally applicable to families of tasks and types of activities rather than being dependent on a particular task. Cognitive models are very specific

representations of part or all aspects of a cognitive theory that apply to a particular task or activity (Gonzalez, 2017). Specifically, normative and descriptive theories of choice often rely on utility theory (Savage, 1954; Morgenstern & Von Neumann, 1953) or aim at describing the psychological impact of perceptions of probability and value on choice (Kahneman & Tversky 1979; 1992). In contrast, models of decisions from experience (DfE) are often dynamic computational representations of sequential choices that are distributed over time and space and that are made under uncertainty (Gonzalez et al., 2017).

Cognitive models of DfE can be used to simulate the interaction of theoretical cognitive processes with the environment, representing a particular task. These models can make predictions regarding how human choices are made in such tasks. These predictions are often compared to data collected from human participants in the same tasks using interactive tools. The explicit comparison of cognitive models' predictions to human actual behavior is a common

---

✉ Cleotilde Gonzalez  
coty@cmu.edu

Thuy Ngoc Nguyen  
ngocnt@cmu.edu

Duy Nhat Phan  
dnphan@andrew.cmu.edu

<sup>1</sup> Carnegie Mellon University, Social and Decision Sciences,  
5000 Forbes Ave., Pittsburgh, 15213 PA, USA

research approach in the cognitive sciences and in particular in the study of decision-making (Gonzalez, 2017).

Cognitive models are dynamic and adaptable computational representations of the cognitive structures and mechanisms involved in decision-making tasks such as DfE tasks under conditions of partial knowledge and uncertainty. Moreover, cognitive models are generative, in the sense that they actually make decisions in similar ways like humans do, based on their own experience, rather than being data-driven and requiring large training sets. In this regard, cognitive models differ from purely statistical approaches, such as machine learning models, that are often capable of evaluating stable, long-term sequential dependencies from existing data but fail to account for the dynamics of human cognition and human adaptation to novel situations.

There are many models of DfE as evidenced by past modeling competitions (Erev et al., 2010; Erev et al., 2017). Most of these models often make broadly disparate assumptions regarding the cognitive processes by which humans make decisions (Erev et al., 2010). For example, the models submitted to these competitions are often applicable to a particular task or choice paradigm rather than presenting an integrated view of how the dynamic choice process from experience is performed by humans. Associative learning models are a class of models of DfE that conceptualize choice as a *learning* process that stores behavior–outcome relationships and are contingent on the environment (Hertwig, 2015). A common example of this type of models is reinforcement learning (RL) (Sutton & Barto, 2018), and the association between DfE and RL is becoming more explicit in the literature (Konstantinidis et al., 2020; Speekenbrink & Konstantinidis, 2015). Generally speaking, these kinds of models rely on learning from reinforcement and the contingencies of the environment as in the Skinnerian tradition (Skinner, 2014; Sutton & Staw, 1995). These models have shown to be successful at representing human learning over time based on feedback.

In contrast to many of the associative learning models, *instance-based learning* (IBL) models rely on a single dynamic decision theory: instance-based learning theory (IBLT) (Gonzalez et al., 2003). IBLT emerged from the need to explain the process of dynamic decision-making, where a sequence of interdependent decisions are made sequentially, over time. IBLT provides a single general algorithm and mathematical formulations of memory retrieval that rely on the well-known ACT-R cognitive architecture (Anderson & Lebiere, 2014). The theory proposes a representation of decisions in the form of *instances*, which are triplets involving state, action, and utilities. In general, states are a representation of the features of the situation of the environment in a task, actions are decisions an agent makes in such states, and utilities are the expectations the agent generates or the outcomes the agent receives

from performing such actions. The theory also provides a process of retrieval of past instances based on their similarity to a current decision situation, and the generation of accumulated value (expectation from experience) based on a mechanism called *blending*, which is a function of the payoffs experienced and the probability of retrieving those instances from memory (Lebiere, 1999; Lejarraga et al., 2012; Gonzalez & Dutt, 2011).

Initially, IBLT was demonstrated in a highly complex, dynamic decision-making task representing the complex process of dynamic allocation of limited resources over time and under time constraints in a “water purification plant” (Gonzalez et al., 2003). Since its inception, many models have been developed based on IBLT, demonstrating human DfE in a large diversity of contexts and domains, from simple and abstract binary choice dynamics (Gonzalez & Dutt, 2011; Lejarraga et al., 2012), to highly specialized tasks such as cyber defense (Aggarwal et al., 2020; Cranford et al., 2020) and anti-phishing detection (Cranford et al., 2019). Also, IBL models have been created to account for dyadic and group effects, where each individual in a group is represented by an IBL agent (Gonzalez et al., 2015). More recently, this IBL algorithm has been applied to multi-state gridworld tasks (Nguyen & Gonzalez, 2020a; 2020b; 2021b) in which the agents execute a sequence of actions with delayed feedback. The recent applications of IBLT have led to significantly more complex and realistic tasks, where multi-dimensional state-action-utility representations are required, where extended training is common, where real-time interactivity between models and humans is needed to solve such tasks (Nguyen & Gonzalez, 2021b).

With the increased use of IBLT in generating models on tasks of greater complexity and in multiple domains, it has become clear that the initial, two-decade-old conceptualization of IBLT needs to be updated. As IBLT has evolved, the initial description of the theory has become less precise, given that no formal implementation of the IBLT process was provided. Thus, a comprehensive description of the current state of the theory along with a concrete implementation of the whole IBL process is essential. Moreover, it is important to demonstrate the full capability and generality of IBLT in a single manuscript, which explains and illustrates how models of multiple and diverse decision tasks can be constructed based on the same theory to generate predictions regarding DfE and learning across a wide range of decision-making tasks. With that, the major goal of this paper is to provide an updated view of the theoretical components of IBLT in a comprehensive and precise form. We also provide an open-source, efficient implementation of the full set of mechanisms of IBLT and demonstrate how such implementation can handle a diverse taxonomy of individual and multi-agent decision-making tasks.

In the process of generating IBL models for more complex tasks that require real-time interactivity between models and humans, we have confronted a practical computational problem, the *curse of exponential growth* (Bellman, 1957; Kuo & Sloan, 2005). The curse of exponential growth is a common problem in models that rely on the accumulation of data over time and on computation of approximate value functions represented as arrays and tables, such as RL models (Sutton & Barto, 2018). As summarized in a recent overview of the challenges in multi-agent RL models, even advanced deep reinforcement learning techniques with many successes in Atari, Go, and StarCraft games (Mnih et al., 2013; Silver et al., 2016; Vinyals et al., 2019) suffer severely from the increase in the dimensions of the state-action space, particularly as the number of agents increases (Wong et al., 2021). The problem becomes even more complex under nonstationary environments and under uncertainty, where information is incomplete. Dynamic conditions significantly increase the diversity and number of states as it is needed for every dynamic decision-making task (Gonzalez et al., 2017). Thus, this paper also addresses the critical question of how IBL models can tackle the curse of exponential growth of memory.

In summary, we present three main contributions. First, an updated view of IBLT provides a comprehensive and precise view of the current theoretical components of the theory, offering a concrete generic algorithm with a formal implementation of the general process of IBLT. Second, we demonstrate the applicability of IBLT across a taxonomy of decision-making tasks varying in the number of agents, the number of actions, the number of decision options and states, and the type of delayed feedback. Third, we provide a new, open-source library, *SpeedyIBL*, that can handle the curse of exponential growth. *SpeedyIBL* allows users to create multiple IBL agents relying on IBLT with fast processing and response time while maintaining the decision characteristics of IBL models. We demonstrate how *SpeedyIBL* is increasingly beneficial (compared to existing implementations, *PyIBL* (Morrison & Gonzalez, 2015)) as the dimensions of the representation, the number of agents and their interactions increase. Through simulation experiments, we demonstrate how IBL models are able to provide predictions across a taxonomy of decision-making tasks with escalating complexity, and how *SpeedyIBL* is increasingly more efficient than *PyIBL* (Morrison & Gonzalez, 2015) as the dimensions of task complexity increase.

## Instance-based learning theory

An updated view of the general decision process proposed in IBLT is illustrated in Fig. 1, and the current mechanisms

of IBLT are made mathematically concrete in Algorithm 1 (Gonzalez et al., 2003).

The process starts with the observation of the environmental state, and the determination of whether there are past experiences in memory (i.e., instances) that are similar to the current environmental state (i.e., recognition). Whether there are similar past instances will determine the process used to generate the expected utility of a decision alternative (i.e., judgment). If there are past experiences that are similar to the current environmental state, the expected utility of such an alternative is calculated via a process of *blending* past instances from memory; but if there are no similar past instances, then the theory suggests that a heuristic is used to generate the expected utility, instead. After judgment, the option with the highest expected utility is maintained in memory and a decision is made as to whether to stop the exploration of additional alternatives and execute the current best decision (i.e., choice) or to continue exploring new alternatives (i.e., exploration loop). When the exploration process ends, the choice that has the highest expected utility is executed, which changes the environment (i.e., execution loop). The loop from recognition to execution continues over time, and the result from a decision may be observed from the environment (i.e., feedback) immediately or with delay from the execution of a choice. Such a decision result (e.g., a reward) is used to update the utility of past instances in memory through a *credit assignment* mechanism.

In IBLT, an “instance” is a memory unit that results from the potential alternatives evaluated. These memory representations consist of three elements which are constructed over time: a situation state  $s$ , which is composed of a set of features  $f$ ; a decision or action  $a$  taken corresponding to an alternative in state  $s$ ; and an expected utility or experienced outcome  $x$  of the action taken in a state.

Each instance in memory has an *activation* value, which represents how readily available that information is in memory, and it is determined by the similarity to past situations, recency, frequency, and noise according to the activation equation in ACT-R (Anderson & Lebiere, 2014). Activation of an instance is used to determine the probability of retrieval of an instance from memory which is a function of its activation relative to the activation of all instances corresponding the same state in memory. The expected utility of a choice option is calculated by blending past outcomes. This blending mechanism for choice has its origins in a more general blending formulation (Lebiere, 1999), but a simplification of this mechanism is often used in models with discrete choice options, defined as the sum of all past experienced outcomes weighted by their probability of retrieval (Gonzalez & Dutt, 2011; Lejarraaga et al., 2012). This formulation of blending represents the general idea of an *expected value* in decision-making, where the probability is a cognitive probability, a function of

the activation equation in ACT-R. Algorithm 1 provides a formal representation of the general IBL process.

**Algorithm 1** Pseudo Code of Instance-based Learning process.

```

Input: default utility  $x_0$ , a memory dictionary  $\mathcal{M} = \{\}$ ,
global counter  $t = 1$ , step limit  $L$ , a flag
 $delayed$  to indicate whether feedback is
delayed.

1 repeat
2   Initialize a counter (i.e., step)  $l = 0$  and observe
   state  $s_l$ 
3   while  $s_l$  is not terminal and  $l < L$  do
4     Execution Loop
5     Exploration Loop  $a \in A$  do
6       Compute activation values  $\Lambda_{i(s_l^j, a)_t}$  of
       instances  $((s_l^j, a), x_{i(s_l^j, a)_t}, T_{i(s_l^j, a)_t})$  by
       Eq. 1
7       Compute retrieval probabilities  $P_{i(s_l^j, a)_t}$ 
       by Eq. 2
8       Compute blended values  $V_{(s_l, a)_t}$ 
       corresponding to  $(s_l, a)$  by Eq. 3
9     end
10    Choose an action  $a_l \in \arg \max_{a \in A} V_{(s_l, a)_t}$ 
11  end
12  Take action  $a_l$ , move to state  $s_{l+1}$ , observe
    $s_{l+1}$ , and receive outcome  $x_{l+1}$ 
13  Store  $t$  into instance corresponding to selecting
    $(s_l, a_l)$  and achieving outcome  $x_{l+1}$  in  $\mathcal{M}$ 
14  If  $delayed$  is true, update outcomes using a
   credit assignment mechanism
15   $l \leftarrow l + 1$  and  $t \leftarrow t + 1$ 
16 end
17 until task stopping condition
    
```

Concretely, for an agent, an option  $k = (s, a)$  is defined by taking action  $a$  after observing state  $s$ .

At time  $t$ , assume that there are  $n_{kt}$  different considered instances  $(k_i, x_{ik_i t})$  for  $i = 1, \dots, n_{kt}$ , associated with  $k$ . Each instance  $i$  in memory has an *Activation* value, which represents how readily available that information is in memory and expressed as follows (Anderson & Lebiere, 2014):

$$\Lambda_{ik_i t} = \ln \left( \sum_{t' \in T_{ik_i t}} (t - t')^{-d} \right) + \alpha \sum_j Sim_j(f_j^k, f_j^{k_i}) + \sigma \ln \frac{1 - \xi_{ik_i t}}{\xi_{ik_i t}}, \tag{1}$$

where  $d$ ,  $\alpha$ , and  $\sigma$  are the decay, mismatch penalty, and noise parameters, respectively, and  $T_{ik_i t} \subset \{0, \dots, t - 1\}$  is the set of the previous timestamps in which the instance

$i$  was observed,  $f_j^k$  is the  $j$ -th attribute of the state  $s$ , and  $Sim_j$  is a similarity function associated with the  $j$ -th attribute. The second term is a partial matching process reflecting the similarity between the current state  $s$  and the state of the option  $k_i$ . The rightmost term represents a noise for capturing individual variation in activation, and  $\xi_{ik_i t}$  is a random number drawn from a uniform distribution  $U(0, 1)$  at each timestep and for each instance and option.

Activation of an instance  $i$  is used to determine the probability of retrieval of an instance from memory.

The probability of an instance  $i$  is defined by a soft-max function as follows

$$P_{ik_i t} = \frac{e^{\Lambda_{ik_i t}/\tau}}{\sum_{j=1}^{n_{kt}} e^{\Lambda_{jk_j t}/\tau}}, \tag{2}$$

where  $\tau$  is the Boltzmann constant (i.e., the “temperature”) in the Boltzmann distribution. For simplicity,  $\tau$  is often defined as a function of the same  $\sigma$  used in the activation equation  $\tau = \sigma \sqrt{2}$ .

The expected utility of option  $k$  is calculated based on *blending* as specified in choice tasks (Lejarraga et al., 2012; Gonzalez & Dutt, 2011):

$$V_{kt} = \sum_{i=1}^{n_{kt}} P_{ik_i t} x_{ik_i t}. \tag{3}$$

The choice rule is to select the option that corresponds to the maximum blended value. In particular, at the  $l$ -th step of an episode, the agent selects the option  $(s_l, a_l)$  with

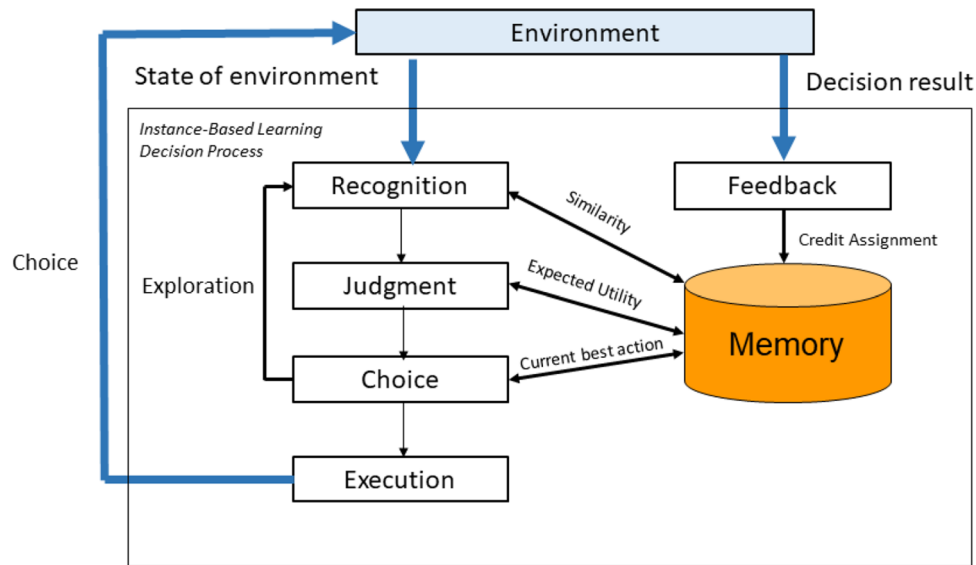
$$a_l = \arg \max_{a \in A} V_{(s_l, a)_t} \tag{4}$$

The flag *delayed* on line 14 of Algorithm 1 is true when the agent knows the real outcome after making a sequence of decision without feedback. In such case, the agent updates outcomes by using one of the credit assignment mechanisms (Nguyen et al., 2021). It is worth noting that when the flag *delayed* is true depends on a specific task. For instance, *delayed* can be set to true when the agent reaches the terminal state, or when the agent receives a positive reward.

**SpeedyIBL implementation**

From the IBL algorithm 1, we observe that its computational cost revolves around the computations on lines 6 (Eq. 1), 7 (Eq. 2), 8 (Eq. 3), and the storage of instances with their associated time stamps on line 13.

Clearly, when the number of states and action variables (dimensions) grow, or the number of IBL agent objects increases, the execution of steps 6 to 8) in Algorithm 1 will directly increase the execution time. The “speedy” version



**Fig. 1** IBLT process from Gonzalez et al. (2003)

of IBL (i.e., SpeedyIBL) is a library focused on dealing with these computations more efficiently.

SpeedyIBL algorithm is the same as that in Algorithm 1. The innovation is in the mathematics. Equations 1, 2 and 3 are replaced with Eqs. 6, 7 and 8, respectively (as explained below). Our idea is to take advantage of *vectorization*, which typically refers to the process of applying a single instruction to a set of values (vector) in parallel, instead of executing a single instruction on a single value at a time. In general, this idea can be implemented in any programming language. We particularly implemented these in Python, since that is how PyIBL is implemented (Morrison & Gonzalez, 2015).

Technically, the memory in an IBL model is stored by using a dictionary  $\mathcal{M}$  that, at time  $t$ , represented as follows:

$$\mathcal{M} = \left\{ k_i : \{x_{ik_{it}} : T_{ik_{it}}, \dots\}, \dots \right\}, \quad (5)$$

where  $(k_i, x_{ik_{it}}, T_{ik_{it}})$  is an instance  $i$  that corresponds to selecting option  $k_i$  and achieving outcome  $x_{ik_{it}}$  with  $T_{ik_{it}}$  being the set of the previous timestamps in which the instance  $i$  is observed.

To vectorize the codes, we convert  $T_{ik_{it}}$  to a NumPy<sup>1</sup> array (Harris et al., 2020) on which we can use standard mathematical functions with built-in Numpy functions for fast operations on entire arrays of data without having to write loops.

After the conversion, we consider  $T_{ik_{it}}$  as a NumPy array. In addition, since we may use a common similarity function for several attributes, we assume that  $f$  is partitioned into  $J$  non-overlapping groups  $f_{[1]}, \dots, f_{[J]}$  with respect to the distinct similarity functions  $Sim_1, \dots, Sim_J$ , i.e.,  $f_{[j]}$

contains attributes that use the same similarity function  $Sim_j$ . We denote  $S(f^k, f^{k_i})$  the second term of Eq. 1 computed by:

```

set  $S(f^k, f^{k_i}) = 0$ 
for  $j = 1$  to  $J$  do
     $S(f^k, f^{k_i}) += \text{sum}((Sim_j(f_{[j]}^k, f_{[j]}^{k_i}))$ 
end for

```

Hence, the activation value (see Eq. 1) can be fast and efficiently computed as follows:

$$\begin{aligned} \Lambda_{ik_{it}} = & \text{math.log}(\text{sum}(\text{pow}(t - T_{ik_{it}}, -d))) \\ & + \alpha * S(f^k, f^{k_i}) \\ & + \sigma * \text{math.log}((1 - \xi_{ik_{it}})/\xi_{ik_{it}}). \end{aligned} \quad (6)$$

With the vectorization, the operation such as `pow` can be performed on multiple elements of the array at once, rather than looping through and executing them one at a time. Similarly, the retrieval probability (see Equation 2) is now computed by:

$$P_{kt} := [P_{1k_{1t}}, \dots, P_{n_{kt}k_{n_{kt}t}}] = v / \text{sum}(v), \quad (7)$$

where  $v = \text{math.exp}([\Lambda_{1k_{1t}}, \dots, \Lambda_{n_{kt}k_{n_{kt}t}t}] / \tau)$ . The blended value (see Equation 3) is now computed by:

$$V_{kt} = \text{sum}(x_{kt} * P_{kt}), \quad (8)$$

where  $x_{kt} := [x_{1k_{1t}}, \dots, x_{n_{kt}k_{n_{kt}t}}]$  is a NumPy array that contains all the outcomes associated with the option  $k$ .

<sup>1</sup><https://numpy.org/doc/stable/>

## Experiments: Demonstration of the general applicability of IBLT

To demonstrate the applicability of IBLT through a wide range of decision tasks as well as to assess the efficiency of SpeedyIBL, we compare SpeedyIBL performance against a regular implementation of the IBL algorithm (Algorithm 1) in Python (PyIBL Morrison & Gonzalez, 2015), in six different tasks that were selected to represent different dimensions of complexity in dynamic decision-making tasks (Gonzalez et al., 2005).

### A taxonomy of individual and multi-agent decision-making tasks

Generally, computational cognitive science has taken advantage of the availability of large amounts of behavioral data to advance the “explanation” of cognitive processes involved in various types of tasks, notably, decision-making (Griffiths, 2015). These models often make excellent predictions of human choices in a particular task. However, for the advancement of cognitive science, it is generally important not to simply make accurate predictions in a specific task but to also provide general explanations and understanding of how and why people behave the way they do across tasks.

The development of computational cognitive models that are based on cognitive theories are expected to provide prediction power without a heavy reliance on data (Hofman et al., 2021). IBLT is a general postulation of mechanisms and processes that are globally applicable to families of dynamic decision tasks, rather than being dependent on the requirements of a particular task. In this section, we present a taxonomy of decision-making tasks that IBLT can address.

Table 1 provides an overview of six dimensions to vary in six different decision-making tasks: (1) number of agents, (2) number of actions, (3) complexity of the states, (4) number of choice options (i.e., alternatives), (5) similarity across states, and (6) feedback delays. The table also presents six tasks that were selected to illustrate how IBLT can handle these dimensions. Although we selected these

six specific tasks to illustrate the generality of IBLT, it is important to note that the theory is applicable to any diversity of tasks within these dimensions. For example IBLT can handle any number of agents, actions, and other task complexities.

In terms of the number of agents, we selected four single agent tasks, one task with two agents, and one task with three agents. The tasks selected for demonstration can have between two to nine potential actions, the number of states and choice options also vary from just a few to a significant large number. We also include one task that requires of similarity judgments across states (i.e., partial matching in Eqs. 1 and 6) and five tasks that do not use similarity judgments. Finally, we include one task with immediate feedback and five tasks that involve feedback delays.

We describe each of the tasks below, starting from the simplest task (repeated *Binary choice*), and moving up in the level of task complexity. The binary choice task has only one state and two options; the *Insider attack* task is a two-stage game in which players choose one of six targets after observing their features to advance. We then scale up to a larger number of states and actions in significantly more complex tasks. A *Minimap* task representing a search and rescue mission and *Ms. Pac-Man* tasks have a larger number of discrete state-action variables. Next, we scale up to two multi-agent tasks: the *Fireman* task has two agents and four actions, and a *Cooperative navigation* task in which three agents navigate and cooperate to accomplish a goal. The number of agents increases the memory computation, since each of those agents adds their own variables to the joint state-action space. Based on these dimensions of increasing complexity, we expect that SpeedyIBL’s benefits over PyIBL will be larger with increasing complexity of the task.

#### Binary choice

In each trial, the agent is required to choose one of two options: Option A or Option B (as illustrated in Fig. 2). A numerical outcome drawn from a distribution after the selection, is the immediate feedback of the task. This is a

**Table 1** Taxonomy of decision-making dimensions, and the illustration of six decision-making tasks

Task	Num. of Agents	Num. of Actions	Num. of States	Num. of Options	Similarity Judgments	Delayed Feedback
Binary choice	1	2	1	2	No	No
Insider attack game	1	6	4	24	Yes	Yes
Minimap	1	4	$\approx 10^{41}$	$\approx 4 \times 10^{41}$	No	Yes
Ms. Pac-Man	1	9	$\approx 10^{347}$	$\approx 9 \times 10^{347}$	No	Yes
Fireman	2	4	$\approx 10^{15}$	$\approx 4 \times 10^{15}$	No	Yes
Cooperative navigation	3	4	$\approx 10^7$	$\approx 4 \times 10^7$	No	Yes

Pick one



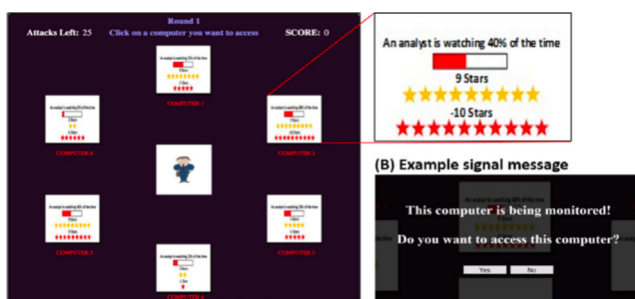
**Fig. 2** Binary choice task

well-studied problem in the literature of risky choice task (Hertwig et al., 2004), where individuals make decisions under uncertainty. Unknown to the agent is that the options A and B are assigned to draw the outcome from a predefined distribution. One option is safe and it yields a fixed medium outcome (i.e., 3) every time it is chosen. The other option is risky, and it yields a high outcome (4) with some probability 0.8, and a low outcome (0) with the complementary probability 0.2.

An IBL model of this task has been created and reported in various past studies, including (Gonzalez & Dutt, 2011; Lejarraga et al., 2012). Here, we conducted the simulations of 1000 runs of 100 trials. We also run the experiment with 5000 trials to more clearly highlight the difference between PyIBL and SpeedyIBL. The default utility  $x_0$  was set to 4.4. For each option  $k$ , where  $k$  is either A or B, we consider all the generated instances taking the form of  $(k, x)$ , where  $x$  is an outcome. The performance is determined by the average proportion of the maximum reward expectation choice (PMax).

### Insider attack game

The insider attack game is an interactive task designed to study the effect of signaling algorithms in cyber deception experiments (e.g., Cranford et al., 2018). Figure 3 illustrates the interface of the task, including a representation of the agent (insider attacker) and the information of six computers. Each of the six computers is “protected” with some probability (designed by a defense algorithm). Each



**Fig. 3** Insider attack game


computer displays the monitoring probability and potential outcomes and the information of the signal. When the agent selects one of the six computers, a signal is presented to the agent (based on the defense signaling strategy); which informs the agent whether the computer is monitored or not. The agent then makes a second decision after the signal: whether to continue or withdraw the attack on the pre-selected computer. If the agent attacks a computer that is monitored, the player loses points, but if the computer is not monitored, the agent wins points. The signals are, therefore, truthful or deceptive. If the agent withdraws the attack, it earns zero points.

In each trial, the agent must decide which of the six computers to attack, and whether to continue or withdraw the attack after receiving a signal. An IBL model of this task has been created and reported in past studies (e.g., Cranford et al. 2019; 2021). We perform the simulations of 1000 runs of 100 episodes. For each option  $(s, a)$ , where the state  $s$  is the features of computers including reward, penalty and the probability that the computers is being monitored (see Cranford et al., 2019 for more details), and  $a \in \{1, \dots, 6\}$  is an index of computers, we consider all the generated instances taking the form of  $(s', a, x)$  with  $s'$  being a state and  $x$  being an outcome. The performance is determined by the average collected reward.

### Search and rescue in Minimap

The Minimap task is inspired by a search and rescue scenario, which involves an agent being placed in a building with multiple rooms and tasked with rescuing victims (Nguyen & Gonzalez, 2021a). Victims have been scattered across the building and their injuries have different degrees of severity with some needing more urgent care than others. In particular, there are 34 victims grouped into two categories (24 green victims and ten yellow victims). There are many obstacles (walls) placed in the path forcing the agent to look for alternative routes. The agent’s goal is to rescue as many victims as possible. The task is simulated as a  $93 \times 50$  grid of cells which represents one floor of this building (see Fig. 4). Each cell is either empty, an obstacle, or a victim. The agent can choose to move left, right, up, or down, and only move one cell at a time.

The agent receives a reward of 0.75 and 0.25 for rescuing a yellow victim and a green victim, respectively. Moving into an obstacle or an empty cell is penalized by 0.05 or 0.01 accordingly. Since the agent might have to make a sequence of decisions to rescue a victim, we update the previous instances by a positive outcome that once the agent receives.



**Fig. 4** Search and rescue map. The empty cells are *white* and the walls are *black*. The *yellow* and *green* cells represent the locations of the yellow and green victims, respectively. The cell with the red color represents the start location of the agent

An IBL model of this task has been created and reported in past studies (Gulati et al., 2021). Here we created the SpeedyIBL implementation of this model to perform the simulation of 100 runs of 100 episodes. An episode terminates when a 2500-trial limit is reached or when the agent successfully rescues all the victims. After each episode, all rescued victims are placed back at the location where they were rescued from and the agent restarts from the pre-defined start position.


In this task, a state  $s$  is represented by a gray-scale image (array) with the same map size. We use the following pixel values to represent the entities in the map:  $s[x][y] = 240$  if the agent locates at the coordinate  $(x, y)$ , 150 if a yellow victim locates at the coordinate  $(x, y)$ , 200 if a green victim locates at the coordinate  $(x, y)$ , 100 if an obstacle locates at the coordinate  $(x, y)$ , and 0 otherwise. For each option  $(s, a)$ , where  $s$  is a state and  $a$  is an action, we consider all the generated instances taking the form of  $(s, a, x)$  with  $x$  being an outcome. The default utility was set to 0.1. The flag *delayed* is set to true if the agent rescues a victim, otherwise false. The performance is determined by the average collected reward.

### Ms. Pac-Man

The next task considered in the experiment is the Ms. Pac-Man game, a benchmark for evaluating agents in machine learning, e.g., Hasselt et al. (2016). The agent maneuvers Pac-Man in a maze while Pac-Man eats the dots (see Fig. 5).

In this particular maze, there are 174 dots and each one is worth 10 points. A level is finished when all dots are eaten.

To make things more difficult, there are also four ghosts in the maze who try to catch Pac-Man, and if they succeed, Pac-Man loses a life. Initially, she has three lives and gets an extra life after reaching 10,000 points. There are four power-up items in the corners of the maze, called power dots (worth 40 points). After Pac-Man eats a power dot, the ghosts turn blue for a short period, they slow down and try to escape from Pac-Man. During this time, Pac-Man is able to eat them, which is worth 200, 400, 800, and 1600 points,



**Fig. 5** Ms. Pac-Man game



consecutively. The point values are reset to 200 each time another power dot is eaten, so the agent would want to eat all four ghosts per power dot. If a ghost is eaten, his remains hurry back to the center of the maze where the ghost is reborn. At certain intervals, a fruit appears near the center of the maze and remains there for a while. Eating this fruit is worth 100 points.

We use the MsPacman-v0 environment developed by Gym OpenAI,<sup>2</sup> where a state is represented by a color image. Here, we developed an IBL model for this task and created the SpeedyIBL implementation of this model to perform the simulation of 100 runs of 100 episodes. An episode terminates when either a 2500-step limit is reached or when Pac-Man successfully eats all the dots or loses three lives. Like in the Minimap task, for each option  $(s, a)$ , where  $s$  is a state and  $a$  is an action, we consider all the generated instances taking the form of  $(s, a, x)$  with  $x$  being an outcome. The parameter *delayed* is set to true if Pac-Man receives a positive reward, otherwise it is set to false. The performance is determined by the average collected reward.

### Fireman

The Fireman task replicates the coordination in firefighting service wherein agents need to pick up matching items for extinguishing fire. This task was used for examining deep reinforcement learning agents (Palmer et al., 2019). In the experiment, the task is simulated in a gridworld of size  $11 \times 14$ , as illustrated in Fig. 6. Two agents A1 and A2 located within the gridworld are tasked with locating an equipment pickup area and choosing one of the firefight items. Afterwards, they need to navigate and find the location of the fire (F) to extinguish it. The task is fully cooperative as both agents are required to extinguish one fire. More importantly, the location of the fire is dynamic in every episode.

The agents receive the collective reward according to the match between their selected firefighting items, which is determined by the payoff matrix in Table 2. The matrix is derived from a partial stochastic climbing game (Matignon et al., 2012) that has a stochastic reward. If they both select the equipment E2, they get 14 points with the probability 0.5, and 0 otherwise. This Fireman task has both stochastic and dynamic properties.

Here we developed an IBL model for this task. We created the SpeedyIBL implementation of this model to




Fig. 6 Fireman game

perform the simulations of 100 runs of 100 episodes. An episode terminates when a 2500-trial limit is reached or when the agents successfully extinguish the fire. After each episode, the fire is replaced in a random location and the agents restart from the pre-defined start positions.

Like in the search and rescue Minimap task, a state  $s$  of the agent A1 (resp. A2) is represented by a gray-scale image with the same gridworld size using the following pixel values to represent the entities in the gridworld:  $s[x][y] = 240$  (resp. 200) if the agent A1 (resp. A2) locates at the coordinate  $(x, y)$ , 55 if the fire locates at the coordinate  $(x, y)$ , 40 if equipment E1 locates at the coordinate  $(x, y)$ , 50 if equipment E2 locates at the coordinate  $(x, y)$ , 60 if equipment E3 locates at the coordinate  $(x, y)$ , 100 if an obstacle locates at the coordinate  $(x, y)$ , 0 otherwise. Moreover, we assume that the agents cannot observe the relative positions of the other, and hence, their states do not include the pixel values of the other agent. For each option

Table 2 Payoff matrix

		Agent 2		
		E1	E2	E3
Agent 1	E1	11	-30	0
	E2	-30	14/0	6
	E3	0	0	5

<sup>2</sup><https://gym.openai.com/envs/MsPacman-v0/>

$(s, a)$ , where  $s$  is a state and  $a$  is an action, we consider all the generated instances taking the form of  $(s, a, x)$  with  $x$  being an outcome. The flag *delayed* is set to true if the agents finish the task, otherwise false. The performance is determined by the average collected reward.

### Cooperative navigation

In this task, three agents (A1, A2, and A3) must cooperate through physical actions to reach a set of three landmarks (L1, L2, and L3) shown in Fig. 7, see Lowe et al. (2017). The agents can observe the relative positions of other agents and landmarks, and are collectively rewarded based on the number of the landmarks that they cover. For instance, if all the agents cover only one landmark L2, they receive one point. By contrast, if they all can cover the three landmarks, they get the maximum of three points. Simply put, the agents want to cover all landmarks, so they need to learn to coordinate the landmark they must cover.

Here we developed an IBL model for this task. We created the SpeedyIBL implementation of this model to perform the simulations of 100 runs of 100 episodes. An episode terminates when a 2500-trial limit is reached or when each of the agents covers one landmark. After each episode, the fire is replaced in a random location and the agents restart from the pre-defined start positions.

In this task, a state  $s$  is also represented by a gray-scale image with the same gridworld size using the following




Fig. 7 Cooperative navigation

pixel values to represent the entities in the environment:  $s[x][y] = 240$  if the agent A1 locates at the coordinate  $(x, y)$ , 200 if the agent A2 locates at the coordinate  $(x, y)$ , 150 if the agent A3 locates at the coordinate  $(x, y)$ , 40 if the landmark L1 locates at the coordinate  $(x, y)$ , 50 if the landmark L2 locates at the coordinate  $(x, y)$ , 60 if the landmark L3 locates at the coordinate  $(x, y)$ , 0 otherwise. For each option  $(s, a)$ , where  $s$  is a state and  $a$  is an action, we consider all the generated instances taking the form of  $(s, a, x)$  with  $x$  being an outcome. The flag *delayed* is set to true if the agents receive a positive reward, otherwise false. The performance is determined by the average collective reward.

### General simulation methods

All the experiments are conducted on a PC AMD 3.00-GHz Ryzen 9 of 16GB RAM and 8 cores with Python 3.7.4 and Numpy 1.19.2. The detailed guideline on how to use the SpeedyIBL package is available at <https://github.com/DDM-Lab/SpeedyIBL> and the Appendix provides a detailed tutorial including installation of the SpeedyIBL library and examples on how to replicate our demonstrations in the tasks offered in this paper.

The parameter values configured in the IBL models with SpeedyIBL and PyIBL implementations were identical. In particular, we used the decay  $d = 0.5$  and noise  $\sigma = 0.25$ . The default utility values generally set to be higher than the maximum value obtained in the task, to create exploration as suggested in Lejarraga et al. (2012) (see the task descriptions for specific values), and they were set the same for PyIBL and SpeedyIBL.

For each of the six tasks, we compared the performance of PyIBL and SpeedyIBL implementations in terms of (i) running time measured in seconds and (ii) performance. The performance measure is identified within each task.

We conducted 1000 runs of the models and each run performed 100 episodes for the *Binary choice* and *Insider attack*. Given the running time required for PyIBL, we only ran 100 runs of 100 episodes for the remaining tasks. We note that an episode of the *Binary choice* and *Insider attack* tasks has one step (trial) while the remaining tasks have 2500 steps within each episode.

The credit assignment mechanisms in IBL are being studied in Nguyen and Gonzalez (2020a). In this paper we used an *equal* credit assignment mechanism for all tasks. This mechanism updates the current outcome for all the actions that took place from the current state to the last state where the agent started or the flag *delayed* was true.

**Table 3** Average running time in seconds of a run

Task	PyIBL time		SpeedyIBL time		Speed up Ratio
Binary choice	0.009		0.008		1.13
Insider attack game	0.141		0.065		2.17
Minimap	21951.88	( $\approx 365$ min $\approx 6$ h)	78.40	( $\approx 1.3$ min)	279.00
Ms. Pac-Man	162372.58	( $\approx 2706.2$ min $\approx 45$ h)	111.98	( $\approx 1.86$ min)	1450.00
Fireman	23743.36	( $\approx 395.72$ min $\approx 6.6$ h)	37.72	( $\approx 0.62$ min)	629.00
Cooperative navigation	9741.37	( $\approx 162$ min $\approx 2.7$ h)	2.59	( $\approx 0.04$ min)	3754.00

## Results

In this section, we present the results of the SpeedyIBL and PyIBL models across all the considered tasks. The comparison these packages is first provided in terms of the average running time and performance, and then in terms of their learning curves.

### Average running time and performance

Table 3 shows the overall average of computational time and Table 4 the average performance across the runs and 100 episodes. The ratio in Table 3 indicates the speed improvement from running the model in SpeedyIBL over PyIBL.

The ratio of PyIBL running time to SpeedyIBL running time in Table 3 shows that the benefit of SpeedyIBL over PyIBL increases significantly with the complexity of the task. In a simple task such as binary choice, SpeedyIBL performs 1.14 faster than PyIBL. However, the speed-up ratio increases with the higher dimensional state space tasks; for example, in Minimap SpeedyIBL was 279 times faster than PyIBL; and in Ms. Pac-Man SpeedyIBL was 1450 times faster than PyIBL.

Furthermore, the multi-agent tasks exhibit the largest ratio benefit of SpeedyIBL over PyIBL. For example, in the cooperative navigation task, PyIBL took about 2.7 h to finish a run, but SpeedyIBL only takes 2.59 s to accomplish a run.

In all tasks, we observe that the computational time of SpeedyIBL is significantly shorter than running the same task in PyIBL; we also observe that there is no significant difference in the performance of SpeedyIBL and PyIBL ( $p > 0.05$ ). These results suggest that SpeedyIBL is able to greatly reduce the execution time of an IBL model without compromising its performance.

### Learning curves

Figure 8 shows the comparison of average running time (middle column) and average performance (right column) between PyIBL (blue) and SpeedyIBL (green) across episodes for all the six tasks.

In the *binary choice* task, it is observed that there is a small difference in the execution time before 100 episodes; where SpeedyIBL performs slightly faster than PyIBL. To illustrate how the benefit of SpeedyIBL over PyIBL implementation increases significantly as the number of

**Table 4** Average performance of a run of 100 episodes

Task	Metric	PyIBL performance	SpeedyIBL performance	<i>t</i> test
Binary choice	PMax	0.833	0.828	$t = -0.83, p = 0.4 > 0.05$
Insider attack game	Average reward	1.383	1.375	$t = -0.38, p = 0.69 > 0.05$
Minimap	Average reward	4.102	4.264	$t = 0.87, p = 0.38 > 0.05$
Ms. Pac-Man	Average reward	228.357	228.464	$t = 0.72, p = 0.47 > 0.05$
Fireman	Average reward	4.783	4.946	$t = 1.07, p = 0.28 > 0.05$
Cooperative navigation	Average reward	2.705	2.726	$t = 0.69, p = 0.48 > 0.05$

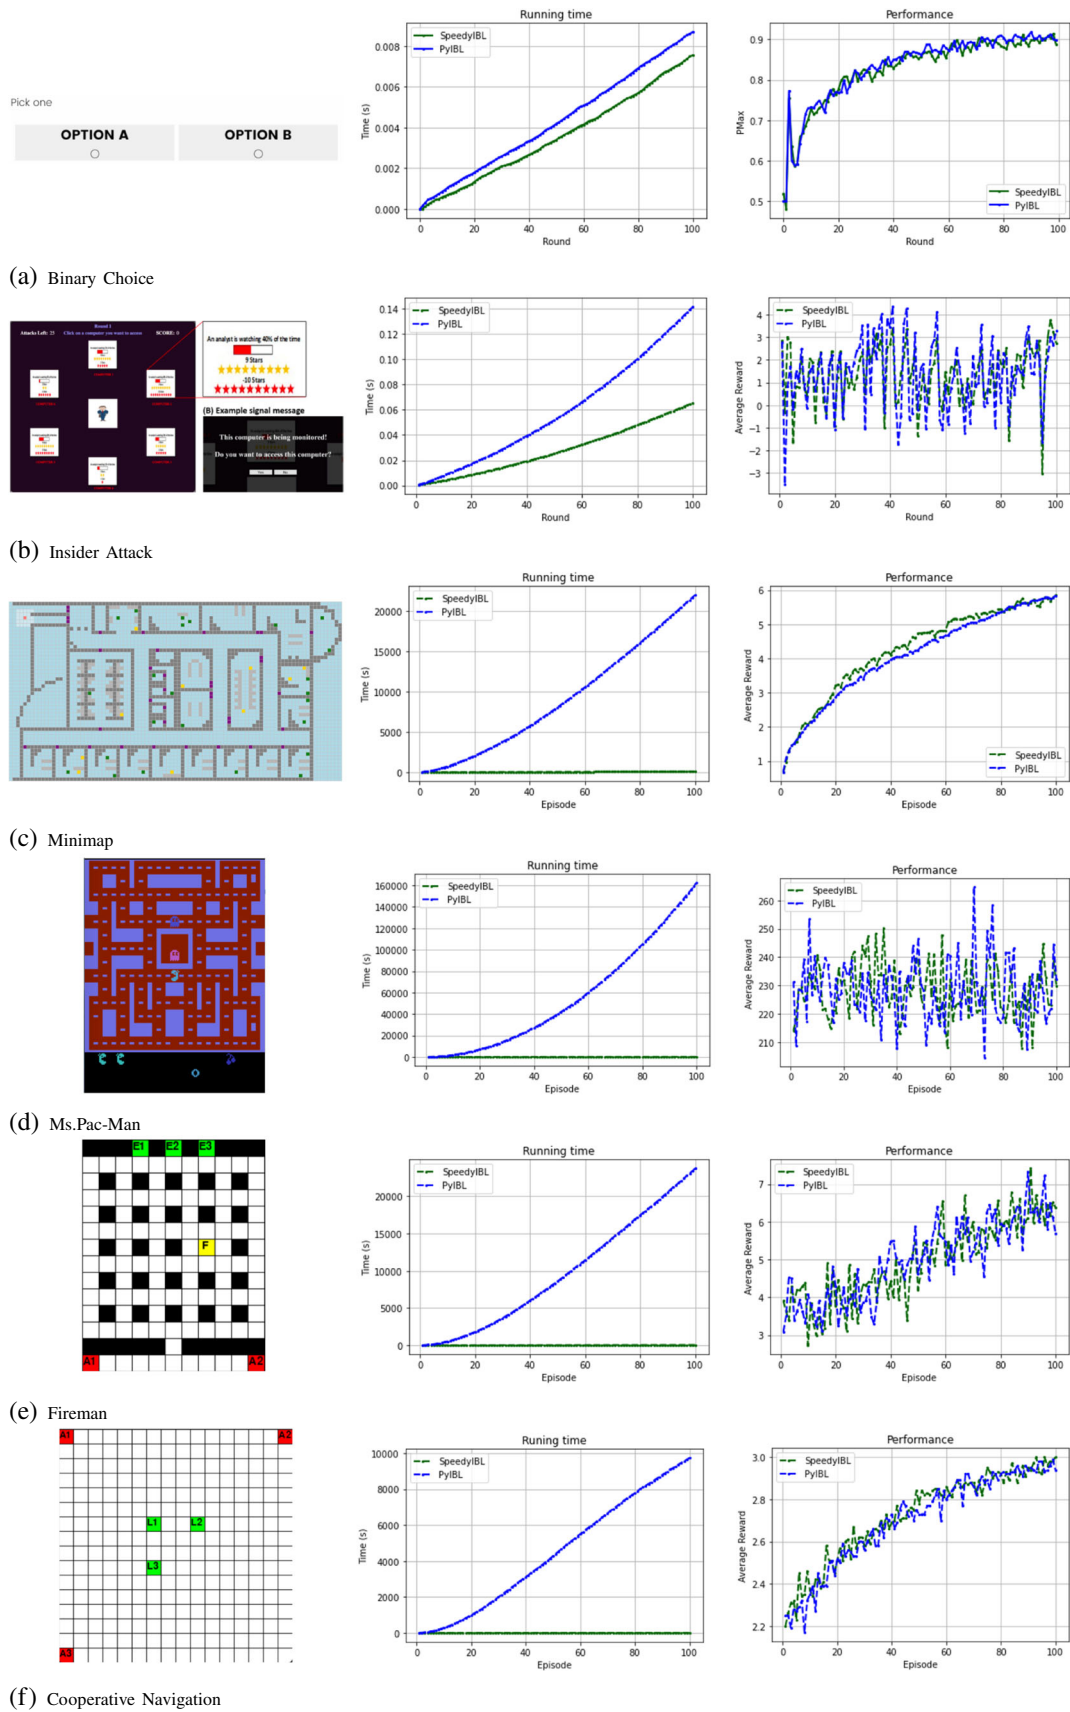



Fig. 8 The comparison between SpeedyIBL (green line) and PyIBL (blue line) over time in the considered tasks



**Fig. 9** The comparison between SpeedyIBL and PyIBL in 5000 playing episodes of binary choice task

episodes increase, we ran these models over 5000 episodes. The results in Fig. 9 illustrate the *curse of exponential growth* very clearly, where PyIBL exponentially increases the execution time with more episodes. The benefit of SpeedyIBL over PyIBL implementation is clear with increased episodes. The PMax of SpeedyIBL and PyIBL overlap, again suggesting no different in their performance.

In the *Insider attack game* as shown Fig. 8a, the relation between SpeedyIBL and PyIBL in terms of computational time shows again, an increased benefit with increased number of episodes. We see that their running time is indistinguishable initially, but then the difference becomes distinct in the last 60 episodes. Regarding the performance (i.e., average reward), again, their performance over time is nearly identical. Learning in this task was more difficult, given the design of this task, and we do not observe a clear upward trend in the learning curve due to the presence of stochastic elements in the task.

In all the rest of the tasks, the *Minimap*, *Ms. Pac-Man*, *Fireman*, and *Cooperative navigation*, given the multi-dimensionality of these tasks representations and the number of agents involved in *Fireman*, and *Cooperative navigation* tasks, the *curse of exponential growth* is observed from early on, as shown in Fig. 8b. The processing time of PyIBL grows nearly exponentially over time in all cases. The curve of SpeedyIBL also increases, but it appears to be constant in relation to the exponential growth of PyIBL given the significant difference between the two, when plotted in the same scale.

The performance over time is again indistinguishable between PyIBL and SpeedyIBL. Depending on the task, the dynamics, and stochastic elements of the task, the models' learning curves appear to fluctuate over time (e.g., *Ms. Pac-Man*), but when the scenarios are consistent over time, the

models show similar learning curves for both, PyIBL and SpeedyIBL.

## Discussion and conclusions

Cognitive models are used increasingly to make predictions of human behavior and simulate the process by which humans make decisions from experience (Cranford et al., 2020; Nguyen & Gonzalez, 2020b; Nguyen et al., 2021). In particular, many computational models have been developed relying on IBLT (Gonzalez et al., 2003). These IBL models have demonstrated how human decision processes are captured and characterized (Gonzalez & Dutt, 2011), and most importantly, they provide evidence for the application and usefulness of the theory.

In this paper, we present an updated account of IBLT, the current formalization of its theoretical components and a comprehensive and precise presentations of the mechanisms of the theory. We aimed at improving the IBLT clarity and describing the mechanisms behind the general process of IBLT with precise mathematical representations and an algorithm implementation. Crucially, we demonstrated the generality and ability of the theory to predict human learning from experience in a wide variety of decision-making tasks. That is, we provided a demonstration of how models grounded on the same IBLT can be applied and handle decision-making tasks varying in the number of agents, the number of actions, the number of decision options and states, and the type of feedback delays.

We observed that implementing IBL models for these tasks using an existing library, PyIBL (Morrison & Gonzalez, 2015), comes at a practical cost. It is difficult to deal with the exponential growth of the memory of instances

as more observations accumulate over time, which leads directly to an exponential slow down of the computational time when the characteristics of the tasks escalate from a single-agent to multi-agent and multi-state settings. Such problem is referred to as the curse of exponential growth, a common computational problem that emerges in many modeling approaches involving tabular computations. Clearly, resolving the curse of exponential growth becomes even more urgent when IBL models are expected to be increasingly used in interactive, real-time tasks that involve humans and models working together, similar to what has been shown recently in a number of RL initiatives (Carroll et al., 2019; Strouse et al., 2021).

To that end, we have developed a new implementation of IBL cognitive models called SpeedyIBL that not only employs a proper data structure for storing memory more efficiently, but also leverages the parallel computation using vectorization (Larsen & Amarasinghe, 2000) to speed up the performance of IBL models in the presence of the curse of exponential growth. We have assessed the robustness of SpeedyIBL by comparing it with PyIBL, a benchmark of the implementation of IBL models in Python (Morrison & Gonzalez, 2015), across a taxonomy of decision-making tasks varying in their increased complexity. We specifically demonstrated that SpeedyIBL implementation is able to perform considerably faster than PyIBL without compromising task performance. Moreover, the results also indicate that the difference in the running time of the SpeedyIBL and PyIBL becomes profound, especially in high-dimensional state spaces and multi-agent domains wherein more agents concurrently collaborate in a task.

Overall, we have introduced SpeedyIBL implementation that enables researchers to create multiple IBL agents relying on IBLT with fast processing and response time. SpeedyIBL can not only be used in simulation experiments of extended learning time, but also can be integrated into browser-based applications in which IBL agents can interact with human subjects in real-time. Given that the computation time of cognitive models in the literature is often overlooked, we believe that the techniques used in SpeedyIBL will be particularly useful for many other ACT-R cognitive models that are still built upon a heavyweight framework programmed in LISP. In that respect, numerous examples can be cited, including a cognitive multi-agent model (Reitter & Lebiere, 2011), a cognitive model for human–robot interaction (Lebiere et al., 2013), hybrid model consisting of a Deep RL agent and a cognitive model (Mitsopoulos et al., 2021), and many other models in the ACT-R literature<sup>3</sup>. Moreover, provided that research

on human–machine behavior has attracted much attention lately, we are convinced that SpeedyIBL will bring significant benefits to researchers and demonstrate the usefulness of IBL models in interactive tasks with human players.

## Transparency and openness

SpeedyIBL is provided as a free and open-source Python library. All the codes, extensive documentation, simulation data, and all scripts used for analyses presented in this manuscript are available on GitHub <https://github.com/DDM-Lab/SpeedyIBL> and on OSF <https://osf.io/gwqte/>. In addition, the Appendix provides a detailed tutorial including installation of the SpeedyIBL library and examples on how to replicate our demonstrations in the tasks offered in this paper.

## Appendix A: SpeedyIBL Tutorial

In an attempt to increase the usage of SpeedyIBL, we hereby provide a tutorial on how to install and use the SpeedyIBL library, following existing research practice (Evans, 2019; Henninger et al., 2021; Vincent, 2016). Specifically, we explain how to build an IBL agent and elaborate on the meaning of associated inputs and functions. Afterwards, we present examples on two illustrative tasks: Binary Choice “Binary choice” and Navigation “Cooperative navigation”. It is worth noting that all the codes to run all the tasks and to reproduce the results presented in the paper are available at <https://github.com/DDM-Lab/SpeedyIBL>. In addition, we provide a Jupyter notebook file of the tutorial, see [https://github.com/DDM-Lab/SpeedyIBL/blob/main/tutorial\\_speedyibl.ipynb](https://github.com/DDM-Lab/SpeedyIBL/blob/main/tutorial_speedyibl.ipynb), for running all the tasks considered in this work using SpeedyIBL. We also make it available on Google Colab [https://colab.research.google.com/github/nhatpd/SpeedyIBL/blob/main/tutorial\\_speedyibl.ipynb](https://colab.research.google.com/github/nhatpd/SpeedyIBL/blob/main/tutorial_speedyibl.ipynb), where one can easily run it with no need to install Python and any relevant modules on their personal computers. Finally, we give a detailed instruction on how to reproduce all the reported results using PyIBL and SpeedyIBL.

## Installing SpeedyIBL

Note that the SpeedyIBL library is a Python module, which is stored at PyPI ([pypi.org](https://pypi.org)), a repository of software for the Python programming language, see <https://pypi.org/project/speedyibl/>. Hence, installing SpeedyIBL is a very simple process. Indeed, one can install SpeedyIBL by simply typing the following line in a command prompt:

<sup>3</sup><http://act-r.psy.cmu.edu/publication/>

```
1 pip install speedyibl
```

## Describing an Agent with SpeedyIBL

After installing the library, we need to import the class `Agent` of SpeedyIBL by typing:

```
1 from speedyibl import Agent
```

We provide the descriptions of the inputs and main functions of the class `Agent` in the following tables.

Inputs	Type	Description
<code>default_utility</code>	float or None	initial utility value for each instance, default = 0.1 or None if prepopulated
<code>noise</code>	float	noise parameter $\sigma$ , default = 0.25
<code>decay</code>	float	decay parameter $d$ , default = 0.5
<code>mismatchPenalty</code>	float or None	mismatch penalty parameter, default = None (without partial matching process)
<code>lendeque</code>	int or None	maximum size of a deque for each instance that contains timestamps or None if unbounded, default = 250000
Functions	Inputs	Description
<code>choose</code>	list of options	choose one option from the given list of options
<code>respond</code>	reward	add the current timestamp to the instance of the last option and reward
<code>prepopulate</code>	option, reward	initialize time 0 for the instance of this option and reward
<code>populate_at</code>	option, reward, time	add time to the instance of this option and reward
<code>equal_delay_feedback</code>	reward, list of instances,	update instances in the list by using this reward
<code>instances</code>	no input	show all the instances in the memory

## Using SpeedyIBL for binary choice task

From the list of inputs of the class `Agent`, although we need five inputs to create an IBL agent, by using the defaults for `noise`, `decay`, `mismatchPenalty`, and `lendeque`, we only need to pass the value of `default_utility` (here in the example is 4.4). Hence we create an IBL agent for the binary choice task as follows:

```
1 agent = Agent(default_utility=4.4)
```

We then define a list of options for the agent to choose:

```
1 options = ['A', 'B'] # A is the safe option while B is the risky one
```

We are now ready to make the agent choose one of the two options:

```
1 choice = agent.choose(options)
```

Next, we determine a reward that the agent can receive after choosing one of the options, see “[Binary choice](#)”:

```

1 import random
2 if choice == 'A':
3     reward = 3
4 elif random.random() <= 0.8:
5     reward = 4
6 else:
7     reward = 0

```

After choosing one option and observing the reward, we use the function `respond`, see the table above, to store the instance in the memory as follows:

```

1 agent.respond(reward)

```

That is, we have run one trial for the binary choice task, which the process includes choosing one option, observing the reward, and storing the instance (`respond`). To conduct 1000 runs of 100 trials, we use two for loops as follows:

```

1 import time # to calculate time
2 runs = 1000 # number of runs (participants)
3 trials = 100 # number of trials (episodes)
4 average_p = [] # to store average of performance (proportion of maximum reward
5     expectation choice)
6 average_time = [] # to save time
7 for r in range(runs):
8     pmax = []
9     ttime = [0]
10    agent.reset() #clear the memory for a new run
11    for i in range(trials):
12        start = time.time()
13        choice = agent.choose(options) # choose one option from the list of two
14        # determine the reward that agent can receive
15        if choice == 'A':
16            reward = 3
17        elif random.random() <= 0.8:
18            reward = 4
19        else:
20            reward = 0
21        # store the instance
22        agent.respond(reward)
23        end = time.time()
24        ttime.append(ttime[-1]+ end - start)
25        pmax.append(choice == 'B')
26    average_p.append(pmax) # save performance of each run
27    average_time.append(ttime) # save time of each run

```



Finally, we provide the following code to plot the running time and performance of this SpeedyIBL agent.

```

1 import matplotlib.pyplot as plt
2 import numpy as np
3 plt.rcParams["figure.figsize"] = (12,4)
4 plt.subplot(int('12'+str(1)))
5 plt.plot(range(trials+1), np.mean(np.asarray(average_time),axis=0), 'o-', color='
    darkgreen', markersize=2, linestyle='--', label='speedyIBL')
6 plt.xlabel('Round')
7 plt.ylabel('Time (s)')
8 plt.title('Runing time')
9 plt.legend()
10 plt.subplot(int('12'+str(2)))
11 plt.plot(range(trials), np.mean(np.asarray(average_p),axis=0), 'o-', color='darkgreen',
    markersize=2, linestyle='--', label='speedyIBL')
12 plt.xlabel('Round')
13 plt.ylabel('PMax')
14 plt.title('Performance')
15 plt.legend()
16 plt.grid(True)
17 plt.show()

```

It is worth noting that the codes of both SpeedyIBL and PyIBL for generating the results of the binary choice task in the paper are available at <https://github.com/DDM-Lab/SpeedyIBL/blob/main/Codes/binarychoice.py>. To plot the results, please see [https://github.com/DDM-Lab/SpeedyIBL/blob/main/Codes/plot\\_results.ipynb](https://github.com/DDM-Lab/SpeedyIBL/blob/main/Codes/plot_results.ipynb).

### Using SpeedyIBL for Cooperative navigation task

First, let us build an environment class of the Cooperative navigation task. Although constructing an environment depends on specific tasks, it consists of two main functions:

`reset` and `step`. The `reset` function sets the agents to their starting locations at beginning of each episode while the `step` function moves the agents to new locations and returns a new state, reward, and task status (task finished or not) after they made decisions.

We would like to note that we created a Python module `vitenv` containing all the environments of the tasks considered in the paper, which can be accessed at <https://pypi.org/project/vitenv/>. The codes of the environments of other tasks and this tutorial also available at our GitHub link <https://github.com/DDM-Lab/SpeedyIBL>. Below is an illustrative code of building the environment of the Cooperative navigation task:

```

1 import numpy as np
2 import copy
3 class Environment(object):
4     def __init__(self):
5         #Initialize elements if the task including size of grid-world, number of agents,
6         #pixel values of agents, landmark, initial locations of agents and landmarks
7         self.GH = 16 #height of grid world
8         self.GW = 16 #width of grid world
9         self.NUMBER_OF_AGENTS = 3 #number of agents
10        self.AGENTS = [240.0, 200.0, 150] #pixel values of [Agent1, Agent2, Agent3]
11        self.LANDMARKS = [40, 50, 60] #pixel values of landmarks
12        self.AGENTS_X = [0, self.GW-1, 0] #x-coordinates of initial locations of agents
13        self.AGENTS_Y = [0, 0, self.GH-1] #y-coordinates of initial locations of agents
14        MID = 8
15        self.LANDMARK_LOCATIONS = [(MID-2,MID-2),(MID+1,MID-2),(MID-2, MID+1)] #locations of
16        #landmarks
17        self.ACTIONS = 4 # move up, down, left, right
18
19    def reset(self):
20        # Reset everything.
21        self.s_t = np.zeros([self.GH,self.GW], dtype=np.float64) #create an array that
22        #represents states of the grid-world
23        # Agents and landmarks are initialised:
24        # Agent x and y positions can be set in the following lists.
25        self.agents_x = copy.deepcopy(self.AGENTS_X)
26        self.agents_y = copy.deepcopy(self.AGENTS_Y)
27        self.agent_status= [False for i in range(self.NUMBER_OF_AGENTS)]
28        #Set pixel values of agents
29        for i in range(self.NUMBER_OF_AGENTS):
30            self.s_t[self.agents_y[i]][self.agents_x[i]] += self.AGENTS[i]
31        #Initialize the landmarks in the environment
32        for l,p in zip(self.LANDMARK_LOCATIONS,self.LANDMARKS):
33            self.s_t[l[1]][l[0]] = p
34        self.reached_landmarks = []
35        return self.s_t
36
37    def step(self, actions):
38        # Change environment state based on actions. :param actions: List of integers
39        #providing actions for each agent
40        #Move agents according to actions.
41        for i in range(self.NUMBER_OF_AGENTS):
42            if not self.agent_status[i]:
43                dx, dy = self.getDelta(actions[i])
44                targetx = self.agents_x[i] + dx
45                targety = self.agents_y[i] + dy

```

```

41     if self.noCollision(targetx, targety):
42         self.s_t[self.agents_y[i]][self.agents_x[i]] -= self.AGENTS[i]
43         self.s_t[targety][targetx] += self.AGENTS[i]
44         self.agents_x[i] = targetx
45         self.agents_y[i] = targety
46         if (targetx,targety) in self.LANDMARK_LOCATIONS:
47             self.agent_status[i] = True
48             if (targetx,targety) not in self.reached_landmarks:
49                 self.reached_landmarks.append((targetx,targety))
50     terminal = sum(self.agent_status) == 3
51     reward = len(self.reached_landmarks)
52     return self.s_t, reward, terminal
53 def getDelta(self, action):
54     # Determine the direction that the agent should take based upon the action selected.
55     # The actions are: 'Up':0, 'Right':1, 'Down':2, 'Left':3, :param action: int
56     if action == 0:
57         return 0, -1
58     elif action == 1:
59         return 1, 0
60     elif action == 2:
61         return 0, 1
62     elif action == 3:
63         return -1, 0
64     elif action == 4:
65         return 0, 0
66 def noCollision(self, x, y):
67     # Checks if x, y coordinate is currently empty :param x: Int, x coordinate :param y:
68     # Int, y coordinate
69     if x < 0 or x >= self.GW or \
70         y < 0 or y >= self.GH or \
71         self.s_t[y][x] in self.AGENTS:
72         return False
73     else:
74         return True

```

Now, we can call the environment and reset it as follows:

```

1 env = Environment()
2 s = env.reset()

```

Like in the binary choice task, we define three agents with `default_utility=2.5` and save them in a list `agents`:

```
1 number_agents = 3
2 agents = []
3 episode_history = {}
4 for i in range(number_agents):
5     agents.append(Agent(default_utility=2.5))
6     episode_history[i] = []
```

Here we have used a dictionary `episode_history` to save information of each episode that we will use for the delay feedback mechanism. Next, we create a list of options:

```
1 s_hash = hash(s.tobytes())
2 options = [(s_hash, a) for a in range(env.ACTIONS)]
```

Here we have used the `hash` function to convert an array into a hashable object used as a key in a Python dictionary. Now we make the agents choose their options and save instances.

```
1 actions = [4,4,4]
2 for i in range(number_agents):
3     if not env.agent_status[i]:
4         option = agents[i].choose(options)
5         actions[i] = option[1]
6         agents[i].respond(0)
7         episode_history[i].append((option[0], option[1], 0, agents[i].t))
```

After choosing actions, the locations of the agents are updated in the environment by the `step` function:

```
1 s, reward, t = env.step(actions)
```

When the agents finish the task (reach landmarks, i.e., `t = True`) or when they reach the maximum number of steps, we update outcomes of previous instances by an equal delayed feedback mechanism.

```
1 for i in range(number_agents):
2     agents[i].equal_delay_feedback(reward, episode_history[i])
```

In order to run 100 times of 100 episodes with 2500 steps, we use the code below.

```

1 import time # to calculate time
2 runs = 100 # number of runs (participants)
3 trials = 100 # number of trials (episodes)
4 steps = 2500 # number of steps
5 average_p = [] # to store average of performance (proportion of maximum reward
   expectation choice)
6 average_time = [] # to save time
7
8 for r in range(runs):
9     preward = []
10    ttime = [0]
11    #clear the memory for a new run
12    for i in range(number_agents):
13        agents[i].reset()
14    episode_history[i] = []
15    for e in range(trials):
16        start = time.time()
17        s = env.reset()
18        # clear the previous episode
19        for i in range(number_agents):
20            episode_history[i] = []
21        for j in range(steps):
22            s_hash = hash(s.tobytes())
23            options = [(s_hash, a) for a in range(env.ACTIONS)]
24            # choose one option from the list
25            actions = [4,4,4]
26            for i in range(number_agents):
27                if not env.agent_status[i]:
28                    option = agents[i].choose(options)
29                    actions[i] = option[1]
30                    agents[i].respond(0)
31                    episode_history[i].append((option[0], option[1], 0, agents[i].t)) # save
information
32            s, reward, t = env.step(actions)
33            if t or j == steps-1:
34                for i in range(number_agents):
35                    agents[i].equal_delay_feedback(reward, episode_history[i])
36                break
37            end = time.time()
38            ttime.append(ttime[-1]+ end - start)
39            preward.append(reward) # save reward of each episode
40            average_p.append(preward) # save performance of each run
41            average_time.append(ttime) # save time of each run

```

To plot the results of the task, we can use the same source code as provided in the binary choice task.

## Reproducing results

All the results can be reproduced by running corresponding scripts for each task under folder **Codes**. In particular, to run the tasks with SpeedyIBL or PyIBL, one can simply execute the following commands and the experiment will start.

### 1. Binary choice task:

```
1 python3 binarychoice.py --method [name]
```

With argument [name] is replaced by: libl for SpeedyIBL and ibl for PyIBL.

### 2. Insider attack game:

```
1 python3 insider_attack_speedyIBL.py # to run SpeedyIBL
2 python3 insider.py # to run PyIBL
```

### 3. Minimap:

```
1 python3 minimap.py --type [name]
```

With argument [name] is replaced by: libl for SpeedyIBL and ibl for PyIBL.

### 4. Ms. Pac-Man:

```
1 python3 mispacman.py --type [name]
```

With argument [name] is replaced by: libl for SpeedyIBL and ibl for PyIBL.

### 5. Fireman:

```
1 python3 fireman.py --type [name]
```

With argument [name] is replaced by: libl for SpeedyIBL and ibl for PyIBL.

### 6. Cooperative navigation:

```
1 python3 navigation.py --type [name]
```

With argument [name] is replaced by: libl for SpeedyIBL and ibl for PyIBL.

**Acknowledgements** This research was partly sponsored by the Defense Advanced Research Projects Agency and was accomplished under Grant Number W911NF-20-1-0006 and by AFRL Award FA8650-20-F-6212 subaward number 1990692 to Cleotilde Gonzalez.

## References

- Aggarwal, P., Thakoor, O., Mate, A., Tambe, M., Cranford, E. A., Lebiere, C., & Gonzalez, C. (2020). An exploratory study of a masking strategy of cyberdeception using CyberVAN. In *Proceedings of the human factors and ergonomics society annual meeting*, (Vol. 64, pp. 446–450). Los Angeles: SAGE Publications Sage CA.
- Anderson, J. R., & Lebiere, C. J. (2014). *The atomic components of thought*. Milton Park: Psychology Press.
- Bellman, R. (1957). *Dynamic programming*. Princeton Univ. New Jersey: Press Princeton.
- Carroll, M., Shah, R., Ho, M. K., Griffiths, T., Seshia, S., Abbeel, P., & Dragan, A. (2019). On the utility of learning about humans for human–AI coordination. *Advances in Neural Information Processing Systems*, 32, 5174–5185.
- Cranford, E. A., Gonzalez, C., Aggarwal, P., Cooney, S., Tambe, M., & Lebiere, C. (2020). Toward personalized deceptive signaling for cyber defense using cognitive models. *Topics in Cognitive Science*, 12(3), 992–1011.
- Cranford, E. A., Gonzalez, C., Aggarwal, P., Tambe, M., Cooney, S., & Lebiere, C. (2021). Towards a cognitive theory of cyber deception. *Cognitive Science* 45(7).
- Cranford, E. A., Lebiere, C., Gonzalez, C., Cooney, S., Vayanos, P., & Tambe, M. (2018). Learning about cyber deception through simulations: Predictions of human decision making with deceptive signals in Stackelberg security games. In Kalish, C., Rau, M. A., Zhu, X. J., & Rogers, T. T. (Eds.) *Proceedings of the 40th annual meeting of the cognitive science society, CogSci 2018, Madison, WI, USA, July 25-28, 2018*.
- Cranford, E. A., Lebiere, C., Rajivan, P., Aggarwal, P., & Gonzalez, C. (2019). Modeling cognitive dynamics in (end)-user response to phishing emails. *Proceedings of the 17th ICCM*.
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions

- under risk, under ambiguity, and from experience. *Psychological Review*, 124(4), 369.
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., . . . , Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47.
- Evans, N. J. (2019). A method, framework, and tutorial for efficiently simulating models of decision-making. *Behavior Research Methods*, 51(5), 2390–2404.
- Gonzalez, C. (2017). Decision-making: a cognitive science perspective. *The Oxford Handbook of Cognitive Science*, 1, 1–27.
- Gonzalez, C., Ben-Asher, N., Martin, J. M., & Dutt, V. (2015). A cognitive model of dynamic cooperation with varied interdependency information. *Cognitive science*, 39(3), 457–495.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating decisions from experience in sampling and repeated choice paradigms. *Psychological Review*, 118(4), 523–51.
- Gonzalez, C., Fakhari, P., & Busemeyer, J. (2017). Dynamic decision making: Learning processes and new research directions. *Human Factors*, 59(5), 713–721.
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27(4), 591–635.
- Gonzalez, C., Vanyukov, P., & Martin, M. K. (2005). The use of microworlds to study dynamic decision making. *Computers in Human Behavior*, 21(2), 273–286.
- Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135, 21–23.
- Gulati, A., Nguyen, T. N., & Gonzalez, C. (2021). Task complexity and performance in individuals and groups without communication. In *AAAI Fall symposium on theory of mind for teams*.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., . . . , Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362.
- Hasselt, H. V., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double q-learning. In *Proceedings of the thirtieth AAAI conference on artificial intelligence, AAAI'16*, (pp. 2094–2100): AAAI Press.
- Henninger, F., Shevchenko, Y., Mertens, U. K., Kieslich, P. J., & Hilbig, B. E. (2021). Lab.js: a free, open, online study builder. *Behavior Research Methods*, 1–18.
- Hertwig, R. (2015). Decisions from experience. *The Wiley Blackwell handbook of judgment and decision making*, 1, 240–267.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539.
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., . . . , et al (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866), 181–188.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, 47(2), 363–391.
- Konstantinidis, E., Harman, J. L., & Gonzalez, C. (2020). Memory patterns for choice adaptation in dynamic environments.
- Kuo, F. Y., & Sloan, I. H. (2005). Lifting the curse of dimensionality. *Notices of the AMS*, 52(11), 1320–1328.
- Larsen, S., & Amarasinghe, S. (2000). Exploiting superword level parallelism with multimedia instruction sets. *ACM SIGPLAN Notices*, 35(5), 145–156.
- Lebiere, C. (1999). Blending: an act-r mechanism for aggregate retrievals. In *Proceedings of the Sixth Annual ACT-R Workshop*.
- Lebiere, C., Jentsch, F., & Ososky, S. Shumaker, R. (Ed.) (2013). *Cognitive models of decision making processes for human-robot interaction*. Berlin: Springer.
- Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: a general model of repeated binary choice. *Journal of Behavioral Decision Making*, 25(2), 143–153.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st international conference on neural information processing systems, NIPS'17*, (pp. 6382–6393). Red Hook: Curran Associates Inc.
- Matignon, L., Laurent, G. J., & Fort-Piat, N. L. (2012). Independent reinforcement learners in cooperative Markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 27(1), 1–31. <https://doi.org/10.1017/S0269888912000057>
- Mitsopoulos, K., Somers, S., Schooler, J., Lebiere, C., Pirolli, P., & Thomson, R. (2021). Toward a psychology of deep reinforcement learning agents using a cognitive architecture. *Topics in Cognitive Science*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing Atari with deep reinforcement learning. arXiv:1312.5602.
- Morgenstern, O., & Von Neumann, J. (1953). *Theory of games and economic behavior*. Princeton: Princeton University Press.
- Morrison, D., & Gonzalez, C. (2015). Pyibl: A python implementation of iblt. <https://www.cmu.edu/dietrich/sds/ddmlab/downloads.html> Version 4.1, Accessed: 2021-09-27.
- Nguyen, T. N., & Gonzalez, C. (2020). Cognitive machine theory of mind. In *Proceedings of CogSci*.
- Nguyen, T. N., & Gonzalez, C. (2020). Effects of decision complexity in goal-seeking gridworlds: a comparison of instance-based learning and reinforcement learning agents. In *Proceedings of the 18th intl conf on cognitive modelling*.
- Nguyen, T. N., & Gonzalez, C. (2021). Minimap: a dynamic decision making interactive tool for search and rescue missions. Tech. rep. Carnegie Mellon University.
- Nguyen, T. N., & Gonzalez, C. (2021). Theory of mind from observation in cognitive models and humans. *Topics in Cognitive Science*. <https://doi.org/10.1111/tops.12553>
- Nguyen, T. N., McDonald, C., & Gonzalez, C. (2021). Credit assignment: Challenges and opportunities in developing human-like ai agents. Tech. rep. Carnegie Mellon University.
- Palmer, G., Savani, R., & Tuyls, K. (2019). Negative update intervals in deep multi-agent reinforcement learning. In Elkind, E., Veloso, M., Agmon, N., & Taylor, M. E. (Eds.) *Proceedings of the 18th international conference on autonomous agents and multiagent systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, (pp. 43–51): International Foundation for Autonomous Agents and Multiagent Systems.
- Reitter, D., & Lebiere, C. (2011). How groups develop a specialized domain vocabulary: a cognitive multi-agent model. *Cognitive Systems Research*, 12(2), 175–185.
- Savage, L. J. (1954). The foundations of statistics. *Naval Research Logistics Quarterly*.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., . . . , et al (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- Skinner, B. F. (2014). *Contingencies of reinforcement: A theoretical analysis*, vol. 3 BF Skinner Foundation.
- Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, 7(2), 351–367.
- Strouse, D., McKee, K., Botvinick, M., Hughes, E., & Everett, R. (2021). Collaborating with humans without human data. *Advances in Neural Information Processing Systems*, 34.

- Sutton, R. I., & Staw, B. M. (1995). What theory is not. *Administrative science quarterly*, 371–384.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: an introduction*. Cambridge: MIT Press.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297–323.
- Vincent, B. T. (2016). Hierarchical Bayesian estimation and hypothesis testing for delay discounting tasks. *Behavior Research Methods*, 48(4), 1608–1620.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., . . . , et al (2019). Alphastar: Mastering the real-time strategy game StarCraft II. *DeepMind blog*, 2.
- Wong, A., Bäck, T., Kononova, A. V., & Plaat, A. (2021). Multiagent deep reinforcement learning: Challenges and directions towards human-like approaches. arXiv:2106.15691.

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.