

A goal-centric outlook on learning

Gaia Molinaro^{1,*} and Anne G. E. Collins^{1, 2}

¹Department of Psychology, University of California, Berkeley, Berkeley, CA, United States

²Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA, United States

*gaiamolinaro@berkeley.edu (G. Molinaro) @gaia_molinaro

June 2023

Abstract

Goals play a central role in human cognition. However, computational theories of learning and decision-making often take goals as given. Here, we review key empirical findings showing that goals shape the representations of inputs, responses, and outcomes, such that setting a goal crucially influences the central aspects of any learning process: states, actions, and rewards. We thus argue that studying goal selection is essential to advance our understanding of learning. By following existing literature in framing goal selection within a hierarchy of decision-making problems, we synthesize important findings on the principles underlying goal value attribution and exploration strategies. Ultimately, we propose that a goal-centric perspective will help develop more complete accounts of learning in both biological and artificial agents.

Keywords goals; learning; decision-making; reinforcement learning; rewards; abstraction; motivation; computational modeling

Highlights

- Recent findings show that goals affect the way in which agents represent their environment, choose relevant actions, and subjectively evaluate outcomes as rewards.
- Because goals influence all aspects of decision-making and learning, developing a complete theory thereof requires that we understand how agents select their goals.
- While computational research on goal generation and management is in its early stages, existing studies suggest that goal selection may be framed as a value-based decision-making process, with values acquired via direct experience or inference, combined with directed and undirected selection policies.

We cannot think, feel, will, or act
without the perception of some goal

Alfred Adler, 1914

Understanding goals to better capture learning and decision-making

In the children’s novel *Matilda* [1], eleven-year-old Bruce Bogtrotter excitedly steals a piece of chocolate cake, and is then forced by headmaster Trunchbull to eat a whole new cake in front of his schoolmates. The task proves grueling, and Bruce quickly feels sick. However, as the audience begins cheering him on, Bruce finishes the cake with revitalized energy. This example, in which the same activity goes from satisfying to nauseating to finally pleasing again over the span of just a few narration acts, hints at the role of internal goals – representations of future states an agent is motivated to realize (Table 1) – in shaping our experiences.

The notion that cognition is vastly goal-dependent dates back to the origins of modern psychology [2], and has been the subject of active research in social, personality, and organizational psychology ever since (e.g., [3–11]). And yet, cognitive scientists have only recently begun to formalize traditional findings and to uncover the neural correlates of goal-dependent representations through neuroimaging and mathematical models [12]. Goals are mostly taken for granted because, in the majority of studies, participants are either trained on or instructed to follow experimenter-defined objectives while individuals’ existing goals tend to be ignored [13]. However, as we argue below, goals shape core aspects of learning and decision-making (see also [14, 15]). Therefore, a complete account of learning is contingent on a better understanding of how goals are selected in the first place.

In the first part of this article, we organize some of the latest research advances on the impact of goals on cognition through the lens of reinforcement learning (RL). First developed from fundamental ideas on animal conditioning and knowledge acquisition, RL has been a fruitful framework for studying how biological and artificial agents learn to make choices [16, 17]. Different formulations of RL share three key components: *states*, which represent the features of the environment the agent is currently considering; *actions*, which the agent can take to change the current state; and *rewards*, which the agent seeks to obtain. While states, actions, and rewards are typically assumed to be fixed and directly available to the agent, each needs to be inferred from context and is heavily dependent on the agent’s goal. For instance, imagine two skiers: one has the goal of finding their friends, while the other one is trying to win a slalom race (Figure 1). Because of these diverging goals, the two skiers will represent the same perceptual input in different ways, consider different sets of actions, and find different outcomes more rewarding than others. In the following paragraphs, we show compelling evidence for the central role of goals in shaping all core aspects of RL, resulting in a redefinition of common terms in the learning literature compared to current standards in social, life, and computer sciences ().

While we use the RL framework to organize relevant findings on goal-dependent cognition, the principles we describe are widely applicable to learning and decision-making; indeed, many of the articles cited below are not part of the RL literature in the traditional sense.

Given the central role of goals in learning, in the second part of this article, we ask what principles might guide goal selection. While this remains a topic of active current research, we begin setting the stage for future work by recasting goal selection as a value-based decision-making problem and drawing insights from studies on curiosity, intrinsic motivation, and autonomous curriculum development.

Goals affect states

Classic RL models assume that states are directly available to learning agents. However, in most situations, appropriate state representations must be learned or inferred [14, 18], often by incorporating latent information [19, 20]. In this section, we review evidence showing that goals affect state representations by guiding the encoding of perceptual inputs and their integration with internal sources of information, sometimes even biasing the processing of perceptual input toward the desired state.

Goal-dependent state representations

Primate cortical regions represent the same perceptual stimuli differently under different goals, showing more similar activation patterns for objects with similar usefulness than stimuli with similar perceptual characteristics but different goal-dependent utility. In an initial study, monkeys were trained to determine either the dominant motion or the color of dots presented on the screen based on a contextual cue. Although the same stimuli were used for the two tasks, neural responses in the prefrontal cortex were contingent on the current task goal [21]. Orthogonal representations for the same stimuli under different task goals have also been found in the human cortex [22]. In another study, human participants imagined using and choosing among certain objects with the goal of either “anchoring a boat” or “starting a fire” while they underwent fMRI scanning. Participants’ choices and neural responses in the various cortical areas reflected the objects’ usefulness under the current goal rather than mere object identity [23].

The principle of goal-dependent state representations may be shared beyond cortical regions, affecting lower-level representations down to early visual areas [24]. Goal-dependent state representations have also been found in hippocampal and parahippocampal regions during both decision-making [25, 26] and navigation tasks [27–29]. Together, these studies confirm the widespread principle of goal-dependent state representations.

Warping perception toward desired states

Goals may even warp internal representations of sensory input toward desirable outcomes, as the longstanding traditions of “motivated perception”, “wishful thinking”, and “optimism bias” suggest [30–33]. In a classic example, Dartmouth and Princeton students watching the same football game believed their own team committed fewer fouls than the opponents [34] – which could not be true for both sides. More recently, [35] provided neurocomputational evidence for the phenomenon of motivated seeing. In their task, participants were required to make decisions (whether an image contained “more face” or “more scene”) on an ambiguous visual input (images that overlapped a face and a scene with varying degrees of transparency). Monetary rewards at the end of the experiment depended on participants’ accuracy, but bonuses were also provided whenever the correct response (regardless of people’s reports) was one of the two categories (e.g., “faces”). Despite the accuracy incentive, participants’ responses and neural representations were biased towards the category they were motivated to see [35], illustrating how goals can pull state representations in the desired direction. While suboptimal in this experiment, goal-dependent representations may benefit learning by facilitating the processing of relevant signals and increasing the motivation to improve (Box 1).

Together, these findings illustrate how, rather than reflecting objective images of externally provided information, state representations are dynamically updated to reflect changes in an agent’s goals. Acknowledging how goals might shape state representations is thus necessary for creating faithful models of the brain and behavior, as well as for developing artificial systems that can operate intelligently in the face of ever-changing needs [36].

Goals affect actions

Goals also exert a direct influence on action representation, selection, and execution – all of which affect the way we study and understand learning.

Action representation

While typical RL algorithms have predefined action spaces, animals need to identify them based on current goals, as meaningful actions might correspond to simple motor responses in some cases, but more abstract actions in others. Motivational constraints can impact which level of abstraction predominates in action representation [37], thus affecting learning. For instance, one study found that people learn stimulus-response associations differently when actions correspond to motor actions (e.g., “press the left-most key”) compared to more abstract responses (“press the key with the blue label, regardless of its position”) [38]. Thus, goals can impact learning by defining what sets of actions should be considered based on a given context. Goals also impact action representations in more applied settings, as supported by the finding that consumer choices while shopping for groceries tend to be organized around related goals [39]. Consistent with these results, the Theory of Event Coding proposes that action representations do not exist in a vacuum, but are rather tightly linked to the current context, including task goals [40].

Goal-dependent alterations in action representation are also reflected at the neuronal level. Different activation patterns have been found in the inferior parietal lobule, an area traditionally involved in motor action representation, for the same grasping movement conducted under different goals (eating vs. placing an object; [41]). Identical actions also display different neural representations in the frontoparietal cortex and lateral-occipital complex based on the goals they serve (e.g., answering “yes” vs. “no” to a question [42]). Therefore, our ability to decode how actions are represented in the brain is limited if goals are not considered.

Action selection

Through their impact on action representation, goals can constrain and direct action selection. This is evident in the finding that goal manipulations can invert the Simon effect – a phenomenon wherein reaction times are affected by the relationship between stimulus and action locations – suggesting that congruency effects are largely dependent on goals, rather than the physical or anatomical location of the effector as previously believed [43]. A more direct role of goals in action selection has been further supported by a set of studies [44] in which participants were instructed to reach either of two equally distant targets in a grid-world task as fast as possible. After planning their trajectory, participants experienced allegedly random disruptions, one of which was designed to bring participants closer to the non-selected target. While the optimal solution would have been to switch targets, participants typically stuck to their original plan, suggesting that active goals restrict how actions are selected.

Goals may also expand the set of available actions based on the task at hand, accounting for seemingly suboptimal action selection strategies. In a series of experiments, children performed tasks efficiently when asked to retrieve specific items in a room, but inserted additional, seemingly unnecessary operations (such as following specific patterns on the floor) if asked to accomplish the same task in a playful way [45]. By considering how internally generated goals may affect individuals’ utility functions, apparently wasteful actions become rational.

Action execution

Goals affect the vigor with which actions are executed, which may in turn impact the speed of learning. In a seminal study, mice were found to increase their speed as they approach the finish line in a corridor

with food at the end [46], suggesting that proximity to a goal energizes motor action in support of goal attainment. Today, energy expenditure in human motion is known to increase as a function of the value assigned to the action’s goal [47–49]. Beyond motor behavior, people also control cognitive effort based on the value associated with attaining a goal [50, 51]. Thus, goals help us predict effort expenditure on both motor and cognitive actions.

Goals affect rewards

In the study of animal neurophysiology and behavior, rewards typically refer to biologically relevant “primary rewards” – e.g., food, water, or sexual pleasure – or “secondary rewards” – stimuli, such as monetary outcomes or numeric points, that have gained reinforcing properties through repeated association with primary rewards. In artificial intelligence, the reward function is typically set by the engineer and designed to solve a specific problem (but see [52]). Viewing rewards as a function of goals expands traditional models, enabling them to capture the richness and flexibility of rewards in complex systems.

Reframing primary and secondary rewards

Although the responses associated with primary rewards are largely innate, intelligent organisms often perceive hedonic properties relative to their goals. For instance, a juicy burger may not look appealing to someone who just had a hefty meal. To account for this phenomenon, homeostatic RL defines primary rewards as a reduction in the discrepancy between current physiological states and homeostatically relevant goals that satisfy, e.g., hunger or thirst [53]. Consistent with this proposal, satiety levels dampen the neural response associated with food [54], which in turn affects the choice strategy [55] and learning performance [56] of individuals working toward the attainment of primary resources. Taking goals into account thus allows for a subjective valuation of primary reward in accordance with internal states.

Given its explanatory power, such a goal-centric perspective on rewards has been extended to include “cognitive setpoints”, which regulate the value attributed to outcomes according to current goals [57, 58]. An illustration of the goal-dependency of rewards is the phenomenon of context-dependent valuation, wherein the subjective value of an outcome depends on the alternative options presented in the same set [59, 60]. The phenomenon can be succinctly explained by recognizing that intrinsically set goals to select the best available option might affect the absolute reward associated with each choice [61].

The role of goals in defining rewards is also evident when the valence of outcomes is determined by personal motivations and task framing. For instance, runners’ satisfaction after a marathon is dependent on initially set time goals [62]. More generally, goals act as reference points in the subjective valuation of outcomes, determining which states should be regarded as gains as opposed to losses [63]. Even external features with strong negative associations may become positively valued if the circumstances make such traits instrumental to goal achievement. Situational goals have been shown to cause preference reversal in animal approach/avoidance responses [64, 65] and human judgments of morality [66]. Task framing manipulations also support the hypothesis that people’s computation of reward is based on goal congruency, rather than value (which is typically conflated with the former): participants’ decision speed for items with higher value tends to be faster in “choose best”, but slower in “choose worst” tasks, compared to items with lower absolute value [67] – a phenomenon likely mediated by goal-dependent attention [68]. Moreover, people regulate their choices so as to maintain a balance across multiple concurrent goals [57], further supporting the idea that people’s value computations are dependent on current goals. Therefore, goals are central to understanding how learning occurs in response to primary and secondary rewards.

Flexible value attribution

The range of outcomes humans and other animals may find rewarding extends beyond primary or secondary reinforcers. In fact, the phrase “intrinsic reward” was coined upon observing that monkeys could spend extended periods of time solving complex puzzles without receiving any extrinsic reward [69]. More recent studies have shown that receiving information elicits similar neural responses as obtaining primary rewards, even if the information itself has no instrumental value [70, 71].

The standard RL theory can easily accommodate intrinsic rewards but remains agnostic regarding the source of intrinsic reward in biological agents [72]. While the cognitive machinery that regulates reward functions has likely been acquired through evolution [73], it is difficult to reconcile the diversity of intrinsic rewards with evolutionary timescales. Indeed, people often spend scarce resources on cognitively taxing or physically effortful activities, ranging from solving Sudoku puzzles to finishing a race, that have little to do with immediate survival [74]. In fact, the value of a stimulus can sometimes be constructed “on the fly” by integrating its observed features with existing knowledge and current goals, providing value estimates before an outcome has even been experienced [75, 76]. Humans can even instantly attribute value to otherwise neutral, novel outcomes and use them as a signal for learning, provided that the outcomes align with current goals [77]. A goal-based reward computation explains how rewards are flexibly implemented in the absence of natural fitness improvements, potentially shedding a better light on the distinction between extrinsic and intrinsic rewards (Box 2). In sum, while rewards are traditionally assumed to be acquired directly from the environment, accounting for goals affords a much richer understanding of where rewards come from and how they are valued [12, 58].

Goal selection as a foundation for learning

We have shown extensive evidence that goals are key in determining how states are represented, actions are selected, and rewards are computed (Figure 1). Given the central role of goals in learning and decision-making, future studies in the field will benefit from a clearer view of subjects’ internal goals as well as the ones imposed by experimental designs [13]. Thus, a computationally precise description of how individuals choose which goals to pursue is warranted. Below, we review and organize existing literature to lay the foundation for how such a critical question may be addressed.

While goal selection acts as a basis for ensuing decisions, some have argued that is not an inherently different problem compared to deciding among subordinate choices [78, 79]. Therefore, the same powerful toolkit the scientific community has developed to study value-based decision-making can also be employed to understand goal selection. Once goal selection is recast as a value-based decision problem, understanding where goals come from reduces to understanding how goal values are set and what policies are employed to select among them, with two interesting caveats. First, goal selection is usually an internal choice that researchers may not directly observe; thus, ingenuity in experimental design or data analysis may be required to investigate goal selection. Second, because of the central role of goals in value learning, the theory can become somewhat circular. However, this exact circularity could explain a range of real-world phenomena, including habits and addiction (Box 3). In the following sections, we examine ways in which goal values have been shown to be estimated through experience, model-based predictions, or proxies based on social cues. We then explore how different types of exploration – directed or undirected – may contribute to the goal selection process.

Experience- and model-based predictions of goal value

Before committing to a goal, people evaluate the costs and benefits associated with its attainment [58, 80–82]. For instance, the goal of becoming a doctor may be set – among other reasons – because

of expectations about the salary associated with the position, and, in turn, the comforts it may afford (Figure 2). The value of some goals is thus tied to innate physiological set points, and may be acquired through direct experience or based on existing world models [83].

Supporting the idea that goal values can be learned, individual performance affects the goal selection process, such that people often reuse and recombine previously pursued goals [84], and tend to return to tasks at which they have succeeded [85–87]. Similarly, people spend more time and effort on goals that have higher value – whether it is acquired directly by observing the outcomes of goal attainment or inferred from knowledge or instructions [51, 88–91].

By repeatedly satisfying primary needs over multiple behavioral trajectories, RL agents can also start developing preferences for specific subgoals – states that, when reached, predict higher success rates for the overall goal – by caching the value associated with setting them [92, 93]. Planning strategies also allow us to set subgoals based on their projected role in satisfying basic needs without direct experience with the subgoals themselves. In this process, memories are generalized and recombined [94] to identify meaningful goals. Therefore, goals can acquire value directly through repeated association with rewards, or indirectly via planning and inference strategies.

Social cues as a proxy for goal value

Social cues can provide a proxy for goal attainability, desirability, and cost, thus circumventing the need for calculating value directly. For instance, one may decide to become a doctor upon observing others pursue similar goals (Figure 2). Indeed, the phenomenon of “goal contagion” – whereby goals spread from one person to the observer – suggests that many of the goals people pursue are derived from others’ [95]. Social cues impact motivation throughout the lifespan and across contexts, supporting people’s decision-making from a young age [96] and affecting their life achievements [97].

Acquiring goals from others relies on the ability to infer others’ objectives from the actions they perform [81, 98]. Such a goal inference process also dictates the level of abstraction at which goals are inherited from others. From a young age, humans can extrapolate higher-order objectives from adult demonstrators’ actions, emulating their goals without necessarily imitating the exact same steps [99], but see [100]. Whether people imitate action sequences or emulate goals often varies based on context, often emulating when goals are easy to infer and resorting to imitation otherwise [101]. In addition, while “minimal” group affiliations (e.g., based on T-shirt color) are often sufficient to induce goal emulation [96], more nuanced and often inferred social characteristics, such as cultural similarity, familiarity, or power dynamics, influence goal contagion [102–104]. Individuals also tend to selectively adopt goals based on how others’ preferences align with their own [105]. Interactions through language also guide which goals people view as available to their own social groups [106]. However, how people adapt the goal contagion process to their own abilities has yet to be specified. Understanding goal inference mechanisms is crucial, because errors in goal inference may give rise to unwanted learning outcomes – a result that is especially relevant for improvements in artificial intelligence, where inverse reinforcement learning (in which reward functions are derived from observed, optimal behaviors [107]) might help align the objectives of artificial systems with human values [108, 109].

Using learning progress to direct goal exploration

Once goal values are set, a policy is required to eventually select one goal. In classic RL, always choosing the current best option could be broadly beneficial, but also lead to overlooking less explored yet potentially better possibilities – a problem known as the exploration-exploitation dilemma [17]. Similar trade-offs might also arise in the context of goal selection, where sampling less known or even out-of-distribution goals could help ensure an adequate estimation of the accessible goal space. Following performance improvement – e.g., choosing to pursue a medical degree because learning

about the human body represents an optimal challenge (Figure 2) – is one possible strategy for agents to discover potentially more valuable goals than ones attempted in the past.

People’s curiosity is often driven towards activities that offer optimal challenges relative to their own capabilities [110]. While humans are generally drawn towards novelty and uncertainty [111], it is an “ideal mismatch” between predictions and observations that seems to drive curious behavior, directing attention towards stimuli of intermediate complexity [112, 113]. Consistent with this, opportunities for control and skill acquisition only feel rewarding when current knowledge is adequate to pursue them [74, 114, 115]. One way to identify optimal levels of uncertainty without imposing arbitrary thresholds is by tracking learning progress, i.e. a measure of change in goal attainment over a series of recent events [116].

Studies of human behavior have confirmed the prominent role of learning progress in dictating which goals people end up pursuing. In choosing among learning tasks with increasing difficulty, people combine internal measures of learning progress with estimates of their own performance, selecting activities that are neither too simple nor impossible to complete [86]. As a result, people naturally tend to pick harder goals as they improve on easier ones [85]. Similarly, when choosing among goals with various levels of uncertainty and perceptual novelty, people favor the ones that maximize opportunities to learn [87], suggesting that people often care more about progress than mere success. People also tend to feel more engaged in tasks with higher mutual information between current goals and means to attain them [117], which benefits learning by directing attention toward action-outcome pairs that are reliable, but not too predictable. Collectively, these studies show that learning progress can successfully guide goal selection toward the acquisition of adaptive skills and away from unachievable feats.

Undirected goal exploration

While learning progress could guide goal exploration strategically, undirected exploration strategies – e.g., entertaining the idea of studying medicine “just because” (Figure 2) – may also contribute to goal selection or even the generation of novel, creative goals. Composing new goals and pursuing far-fetched ones creates opportunities to explore the environment in a structured manner, which ultimately improves performance while attempting more demonstrably useful goals [84, 116, 118].

Children master goal generation during play. Indeed, games can be defined as a voluntary submission to arbitrary goals and rules associated with their accomplishment [119]. While some of the goals children come up with during play (like skipping a rope or winning a race) may help train practical skills, even seemingly meaningless goals (such as building a rocket ship or fighting a dragon) could be beneficial as they bring about opportunities to structure thoughts, innovate, and learn [45]. By pursuing a goal – any goal, regardless of the outcome it yields – children learn how to flexibly assign value to otherwise neutral states of the world, much like what is required of adults to accomplish more instrumental feats [77].

The process of exploring creative goals and pursuing new ones without serious intentions may continue during adulthood, with mind wandering and daydreaming replacing the physical engagement towards bizarre goals, although it is also possible that goal setting becomes more “directed” with age [120]. Admittedly, however, we still know very little about how inventive goals – whether realistic or utterly impossible – support the acquisition of useful skills.

Concluding Remarks

A growing literature shows how neural and behavioral correlates of learning – and, likely, other cognitive processes – could be better understood by recognizing the central role of goals. On the model of Dobzhansky’s famous quote about biology and evolution [121], one could argue that nothing in

psychology makes sense except in the light of goals. Nonetheless, comprehensive models of how goals are created, set, and pursued are currently lacking. Recasting the goal selection process as a higher-level decision-making problem suggests that goals may be chosen based on the value attributed to their attainment, which may be acquired through experience, predicted based on prior knowledge, or proxied through heuristics such as following others' lead. Goal selection policies could then maximize goal value, or direct exploration toward alternative goals. Providing computationally precise theories of goal selection, as well as goal maintenance, prioritization, and pursuit, is a key challenge for future research, alongside a more thorough appreciation of how goals and other cognitive components of learning interact (see Outstanding Questions). Particularly interesting work could be done in attempting to bridge research from social, personality, and organizational psychology with computational frameworks from cognitive and computer science – especially in addressing how goals at different hierarchical levels of representation may be controlled. Because of the acknowledged importance of language in learning (e.g., [122]), future studies may also investigate how language supports the creation of new goals and the maintenance of existing ones [123]. Addressing such questions will likely require the development of better tools to infer participants' intrinsic goals and novel experimental paradigms that treat goal selection as an outcome of interest, rather than an independent variable or a task feature. Ultimately, more precise frameworks of how goals originate and get selected will lead to major advances in our understanding of how individuals learn, often autonomously, in complex worlds.

Table 1

	Psychology/Neuroscience	Computational RL	This paper
Action	“A self-initiated sequence of movements, usually with respect to some goal” ⁱ	“The mechanism by which the agent transitions between states of the environment” ⁱⁱ	A means by which an agent interacts with, and potentially manipulates, its environment
Environment	“The aggregate of external agents or conditions — physical, biological, social, and cultural — that influence the functions of an organism” ⁱ	The external and internal system with which an agent interacts and receives feedback from ⁱⁱ [17]	The set of circumstances (external and internal) in which an agent operates and from which it receives input
Goal	“The end state toward which a human or nonhuman animal is striving: the purpose of an activity or endeavor” ⁱ	The objective that the agent optimizes [17]	A representation of a future state an agent is motivated to realize [2]
Motivation	“A person’s willingness to exert physical or mental effort in pursuit of a goal or outcome” ⁱ	NA	The physiological and psychological driving force that initiates, energizes, and sustains goal-directed behavior
Reward	A positive stimulus or event that an organism experiences or approaches, which leads to an increase in the likelihood of a particular behavior or action being repeated in the future [124]	“The numerical result of taking an action in a state, as defined by the environment” ⁱⁱ , which an agent seeks to maximize [73]	A signal that measures the agent’s performance with respect to its current goal
State	“The condition or status of an entity or system at a particular time that is characterized by relative stability of its basic components or elements” ⁱ	The feature values that “describe the current configuration of the environment, which the agent uses to choose an action.” ⁱⁱ	A configuration of the (internal or external) environment as represented by an agent, who uses it to inform behavior
Task	“Any goal-oriented activity undertaken by an individual or a group” ⁱ	A problem a system is designed to solve, defined by a set of inputs, transition functions, and output characteristics [125]	An assigned activity or problem to solve
Value	“The worth, usefulness, or importance attached to something” ⁱ	The expected long-term return with discount, as opposed to immediate reward [17]	An agent’s estimate of the potential long-term benefit it can expect from a state
Vigor	“Physical and mental robustness and energy” ⁱ	NA	The intensity with which an agent performs an action or the effort it exerts in attempting it

Table 1: **Learning-related terms across fields.** The same terms have been traditionally used in psychology/neuroscience and artificial intelligence (particularly the computational reinforcement learning (RL) subfield) to describe slightly different constructs. Recognizing the prominent role of goals in learning warrants updating the meaning of some of these concepts when referring to biological agents, which may in turn prompt innovation in the field of artificial intelligence. Note that some terms do not have applicable definitions in artificial intelligence. While most definitions are taken or adapted from generally accepted sources, such as the APA Psychology Dictionaryⁱ and the standard textbook for computational RL [17], different sub-disciplines within psychology and neuroscience or computational RL adopt these terms in even more nuanced ways. Our intent here is primarily to show how goals may be integrated into current definitions.

Figure 1

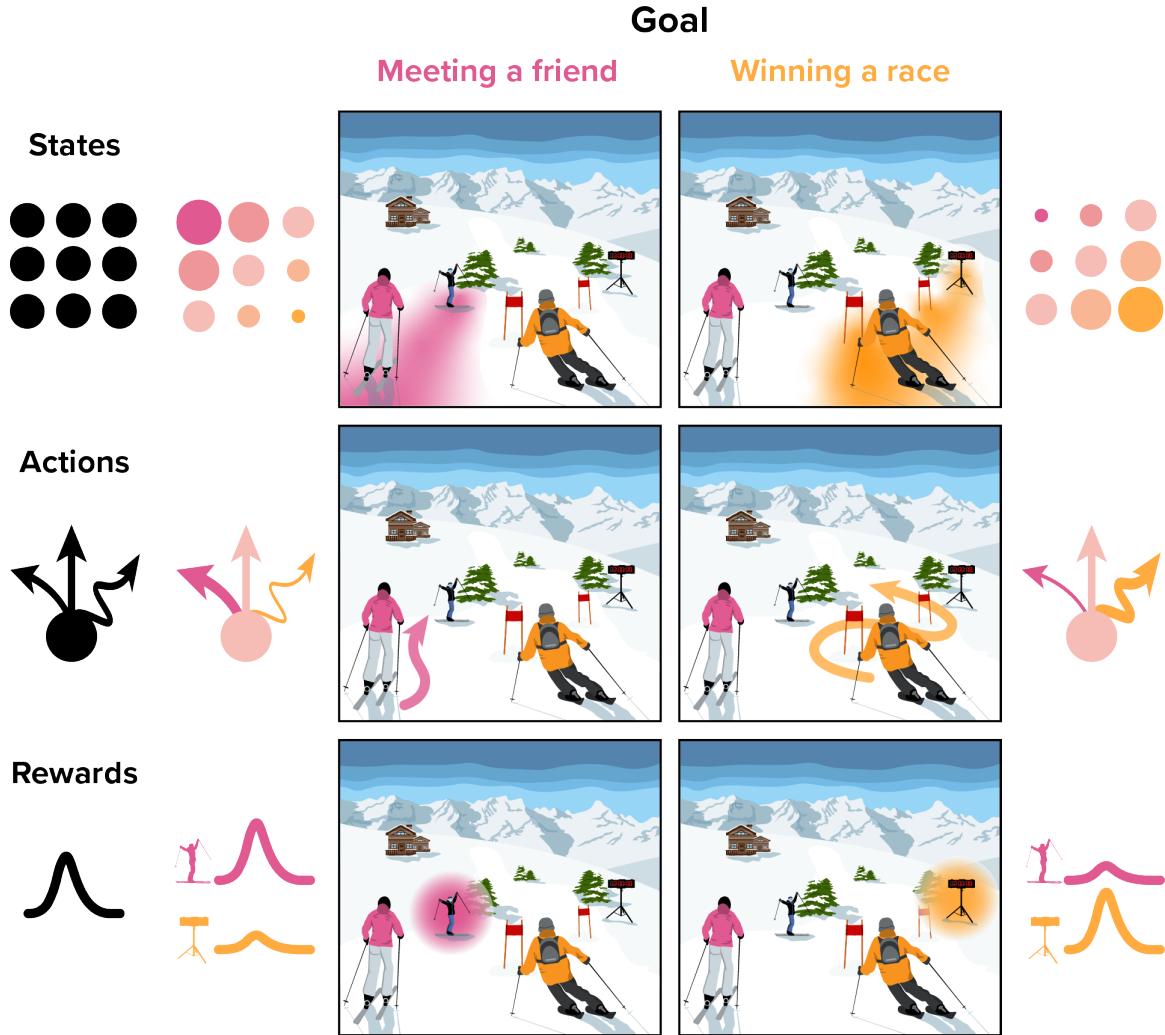


Figure 1: **Goals impact key components of learning: states, actions, and rewards.** Two skiers are presented with similar perceptual input, but have the separate goals of meeting a friend (left, in pink) and winning a slalom race (right, in orange). **Top row:** Having different goals impacts the way in which states are represented by enhancing certain aspects of incoming sensory information over others. **Middle row:** Goals impact action representation, selection, and execution by affecting which family of actions is considered and eventually chosen (e.g., actions that lead to a destination with the shortest possible path vs. actions that pass slalom gates in the correct direction) as well as, potentially, the vigor with which it is carried out. **Bottom row:** Goals determine which states trigger reward signals (e.g., reaching one's friend's location vs. finishing the race in a short time).

Figure 2

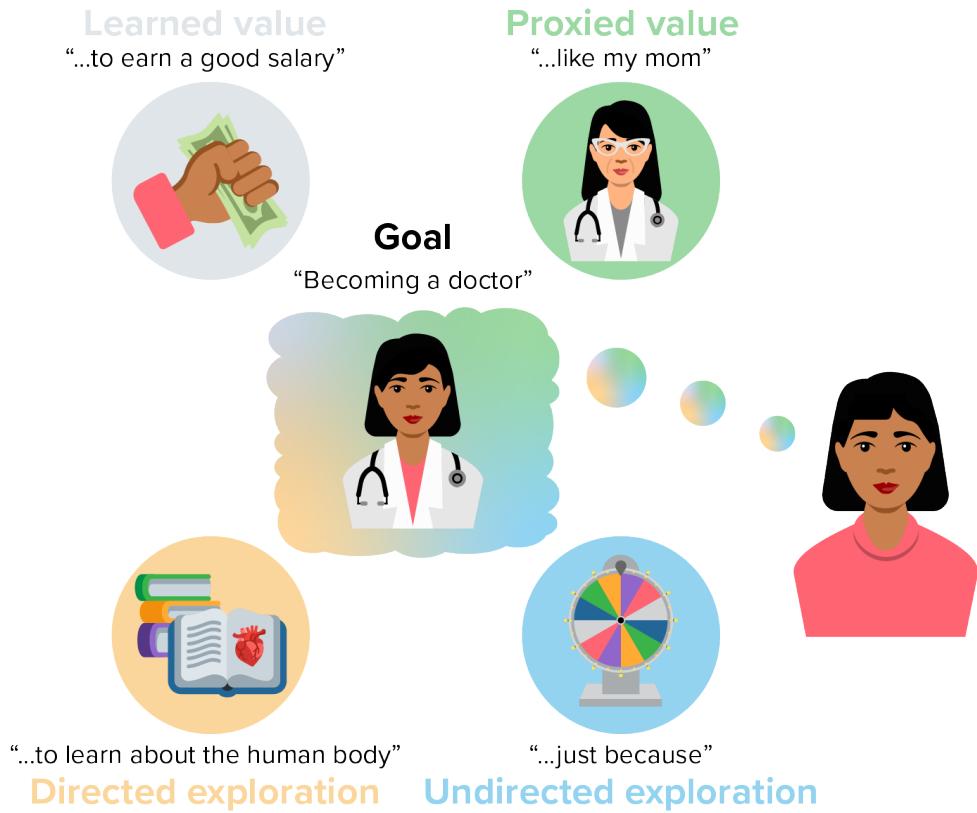


Figure 2: Value abstraction, social cues, and self-monitoring impact goal selection. In the example above, the goal of becoming a doctor could originate from (any, or a mixture of) various elements. On the one hand, selecting a specific goal depends on the value attributed to it, which may be based on experience or one's models of the world (predicting, e.g., a good salary as compensation), or proxied by observing others (e.g., a parent) pursue similar goals. On the other hand, goal selection policies may account for exploration strategies, guiding goal selection based on learning progress (e.g., measured as positive changes in one's knowledge about the human body), or even choosing goals in an undirected fashion.

Box 1

The pros and cons of goal-dependent representations

Compared to absolute information coding, goal-contingent representations are potentially adaptive, as they minimize interference in the selection of actions that best serve the current goal [24]. Naturalistic stimuli tend to be highly complex, and thus unsuitable for the computation of adequate responses within useful timescales. By directing the focus of attention towards relevant stimuli, goal-dependent representations can significantly simplify information processing according to current needs [126]. A similar principle applies to the action space, as restricting the field to goal-relevant actions can significantly ease the action selection process [127]. The goal-dependent attribution of rewarding properties to otherwise neutral outcomes is especially useful for flexible learning, as it allows intrinsically motivated agents (both human and artificial) to acquire skills adaptively in the absence of external incentives [52]. Biasing perception towards desired states may also benefit learning: in fact, people tend to perform better in learning tasks when they believe – even if wrongly – they are approaching goal achievement [128–131]. For some, goals are even valuable in and of themselves, as they promote the generation of new ideas to serve adaptive behavior as a whole [132].

However, reconfiguring neural processing systems based on one’s goal is costly. As a result, people are often resistant to switching goals. The classic “sunk cost fallacy” phenomenon, as well as recent experimental findings, show that people often keep pursuing goals they have previously adopted even when they are no longer beneficial [44, 132] – potentially due to switch costs [133].

A potential solution for maximizing the advantages of goal-dependent representations in light of their costs was suggested by the “Rubicon model” and its later adaptations. According to this view, prior to committing to a goal, people evaluate the costs and benefits associated with it and hold a conservative perspective. Once a goal has been selected (and the “Rubicon has been crossed”), entire sets of representations are reconfigured to funnel information processing towards goal completion, and people become optimistic about their ability to achieve the goal [58, 80].

Box 2

On distinguishing intrinsic and extrinsic rewards

Distinguishing between intrinsic and extrinsic rewards is deemed important for practical applications, e.g., in education or public policy, because the quality of learning often varies based on people's motivations [115]. Understanding motivations is also of central interest for research with human and non-human animal subjects, who are often tasked with performing cognitive or behavioral tasks for the sake of some extrinsic reward (e.g., juice, points, money, course credit). However, identifying the exact boundary between intrinsic and extrinsic rewards has been challenging, and defining rewards as a function of specific goals may provide a more efficient approach.

In the classic definition, intrinsic motivation refers to the willingness to engage in an activity because it is inherently enjoyable and often dissociated from immediate survival benefits [74, 134], while extrinsic motivation refers to the willingness to take part in an activity because of the future separate outcomes it is expected to yield [115]. Intrinsic motivation has also been characterized by its dissociation from immediate survival benefits and ties to self-efficacy, i.e., one's belief that one is capable of attaining immediate or future goals [74].

Nonetheless, some types of rewards, such as social rewards or physical fitness, are difficult to categorize as either intrinsic or extrinsic. Moreover, people who are seemingly intrinsically motivated to perform an activity for the sake of it (e.g., an amateur painter) often concurrently aim for specific outcomes (a beautiful painting) [135]. Extrinsic rewards may also be internalized into one's own value system, posing additional challenges to the distinction between intrinsic and extrinsic rewards [115]. Even in the artificial intelligence literature (where intrinsic rewards are computed by the agent itself while extrinsic rewards are provided by a predefined optimization function), the distinction fades when considering machines that manipulate reward functions by setting their own goals [116, 136].

These caveats suggest that rather than being characterized by inherently distinct qualities, intrinsic and extrinsic rewards lie on a spectrum. For researchers involved in animal studies, it would perhaps be more useful to recognize goals the subjects may generate internally alongside goals defined by the experimenter [13, 61] while remaining agnostic on whether they should be classified as intrinsic or extrinsic.

Box 3

Interactions between learning and goal values

As highlighted in the first half of this article, goals have a central role in the way states, actions, and rewards are represented, controlled, and delivered. But if we cast goal selection as a value-based decision-making process – as proposed in the second part of the review – the theory becomes circular [12] since, in turn, goal setting depends on how we represent, choose, and value the goals themselves. While the recursiveness of this theory makes it harder to study, it is likely integral to the cognitive processes involved in learning and decision-making, and may even explain various real-life phenomena.

The fact that goals shape the representation of current states (but also the retrieval of information from the past [137]), suggests that, in turn, the space of potential goals may be biased by current ones. Hence, people will be more likely to select goals akin to previously chosen ones – which may explain how specific interests develop over time.

When choosing which goal to pursue, people consider the costs and feasibility of the actions required to attain it [58, 79]. Thus, the role of goals in shaping the action space eventually reflects back to the goal selection process itself.

Because goals affect the way in which rewards are calculated, obtaining rewards could, in turn, increase the cached value of goals themselves – making them more likely to be selected again in the future. Indeed, people tend to return to tasks that previously generated rewarding outcomes [85–87]. Moreover, goal selection can be learned in a manner analogous to action selection, which may underlie the formation of habits [78, 118, 138, 139]. Similarly, compulsive drug-seeking can be explained in terms of a goal selection pathology, wherein drug abuse causes reward boosts which, in turn, lead to a persistent pursuit of goals pertaining to the drug [140].

Given the complex interactions reported above, the relationship between goal-dependent representations and the value of goals themselves should be the object of future study. Nonetheless, the issue does not negate the fact that goals have a profound impact on learning, and should therefore be addressed in order to fully understand how organisms and machines acquire intelligent behavior.

Outstanding Questions

- What neural and cognitive mechanisms allow goals to shape state and action spaces?
- How does the brain translate reaching a goal into a reward signal? While executive functions and learning processes are likely to be involved, the exact mechanisms of how goals may acquire rewarding properties remain unclear.
- How do learning and goal setting processes interact?
- What other factors contribute to goal selection?
- How do different factors jointly contribute to goal selection? Previous research suggests various elements interact non-linearly in the setting and maintenance of goals, but accurate precise descriptions of such relationships are currently missing.
- How are multiple goals monitored, managed, and prioritized, especially in the presence of conflict?
- How can ideas on goal selection from social and personality psychology be formalized by computational theories?
- How do people choose appropriate abstraction levels for goal selection, and how do goals at different hierarchical levels interact with each other?
- How can we formalize the recursive relationship between goal selection and reinforcement learning?
- How can “tasks” be distinguished from “goals” in computational models?
- How does goal selection change across development?
- What factors characterize the distinction between explicit cognitive goals and implicit or homeostatic ones?
- What tools may researchers use to study goal selection as a dependent variable? The development of new strategies to measure voluntarily chosen goals will require ingenious experimental designs and novel techniques to tap into participants’ internal motivations.
- How does language support the creation of new goals and the maintenance of existing ones? Given its compositionality, language is an invaluable tool for the invention of original goals. Moreover, the way in which goals are expressed in language – spoken or thought – likely plays an important role in how those goals impact cognition.

Acknowledgments

We are grateful to all Computational Cognitive Neuroscience Lab (University of California, Berkeley) members for providing insightful feedback on an early draft of the manuscript.

Resources

- i <https://dictionary.apa.org/>
- ii <https://developers.google.com/machine-learning/glossary/r/>

References

- [1] Dahl, R. (1988) *Matilda*. Jonathan Cape
- [2] Elliot, A.J. and Fryer, J.W. (2008) The goal construct in psychology. In *Handbook of motivation science* (Shah, Y.J. and Gardner, W.L., eds.), pp. 235–250. The Guilford Press
- [3] Carver, C.S. and Scheier, M.F. (1998) *On the Self-Regulation of Behavior*. Cambridge University Press
- [4] Gollwitzer, P.M. (2018) The goal concept: A helpful tool for theory development and testing in motivation science. *Motiv. Sci.* 4, 185–205
- [5] Höchli, B. et al. (2019) Making new year's resolutions that stick: Exploring how superordinate and subordinate goals motivate goal pursuit. *Appl. Psychol. Health Well-Being* 12, 30–52
- [6] Jia, L. et al. (2019) Protecting an important goal: When prior self-control increases motivation for active goal pursuit. *J. Exp. Soc. Psychol.* 85, 103875
- [7] Kruglanski, A.W. et al. (2002) A theory of goal systems. In *Advances in Experimental Social Psychology*, pp. 331–378. Elsevier
- [8] Kruglanski, A.W. et al., eds. (2023) *Goal Systems Theory*. Oxford University Press New York
- [9] Kung, F.Y.H. and Scholer, A.A. (2021) Moving beyond two goals: An integrative review and framework for the study of multiple goals. *Personality and Social Psychology Review* 25, 130–158
- [10] Locke, E.A. and Latham, G.P. (2002) Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *Am. Psychol.* 57, 705–717
- [11] Sheldon, K.M. et al. (2019) Rightly crossing the Rubicon: Evaluating goal self-concordance prior to selection helps people choose more intrinsic goals. *J. Res. Pers.* 79, 119–129
- [12] Juechems, K. and Summerfield, C. (2019) Where does value come from? *Trends Cogn. Sci.* 23, 836–850
- [13] Karayanni, M. and Nelken, I. (2022) Extrinsic rewards, intrinsic rewards, and non-optimal behavior. *J. Comput. Neurosci.* pp. s10827–022–00813–z
- [14] De Martino, B. and Cortese, A. (2022) Goals, usefulness and abstraction in value-based choice. *Trends Cogn. Sci.*
- [15] Frömer, R. and Shenhav, A. (2021) Filling the gaps: Cognitive control as a critical lens for understanding mechanisms of value-based decision-making. *Neurosci. Biobehav. Rev.*
- [16] Niv, Y. (2009) Reinforcement learning in the brain. *J. Math. Psychol.* 53, 139–154

- [17] Sutton, R.S. and Barto, A.G. (2018) *Reinforcement Learning: An Introduction*. MIT Press
- [18] Gershman, S.J. and Niv, Y. (2010) Learning latent structure: carving nature at its joints. *Curr. Opin. Neurobiol.* 20, 251–256
- [19] Schuck, N.W. et al. (2016) Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* 91, 1402–1412
- [20] Wilson, R. and Niv, Y. (2012) Inferring relevance in a changing world. *Front. Hum. Neurosci.* 5
- [21] Mante, V. et al. (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84
- [22] Flesch, T. et al. (2022) Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* 110, 1258–1270.e11
- [23] Castegnetti, G. et al. (2021) How usefulness shapes neural representations during goal-directed behavior. *Sci. Adv.* 7, eabd5363
- [24] Schaffner, J. et al. (2023) Sensory perception relies on fitness-maximizing codes. *Nat. Hum. Behav.* pp. 1–17
- [25] Edelson, M.G. and Hare, T.A. (2021). Goal-dependent hippocampal representations facilitate self-control. *bioRxiv*. Published online August 28, 2021. <https://doi.org/10.1101/2021.08.26.457750>
- [26] Park, S.A. et al. (2020) Map making: constructing, combining, and inferring on abstract cognitive maps. *Neuron* 107, 1226–1238
- [27] Crivelli-Decker, J. et al. (2023) Goal-oriented representations in the human hippocampus during planning and navigation. *Nat. Commun.* 14, 2946
- [28] Muhle-Karbe, P.S. et al. (2023). Goal-seeking compresses neural codes for space in the human hippocampus and orbitofrontal cortex. *bioRxiv*. Published online February 3, 2023. <https://doi.org/10.1101/2023.01.12.523762>
- [29] Nyberg, N. et al. (2022) Spatial goal coding in the hippocampal formation. *Neuron*
- [30] Balcetis, E. and Dunning, D. (2006) See what you want to see: motivational influences on visual perception. *J. Pers. Soc. Psychol.* 91, 612–625
- [31] Kruglanski, A.W. et al. (2020) All thinking is ‘wishful’ thinking. *Treds Cogn. Sci.* 24, 413–424
- [32] Sharot, T. (2011) The optimism bias. *Curr. Biol.* 21, R941–R945
- [33] Sharot, T. (2019) Is visual representation coloured by desire? *Nat. Hum. Behav.* 3, 891–892
- [34] Hastorf, A.H. and Cantril, H. (1954) They saw a game; a case study. *J. Abnorm. Soc. Psychol.* 49, 129–134
- [35] Leong, Y.C. et al. (2019) Neurocomputational mechanisms underlying motivated seeing. *Nat. Hum. Behav.* 3, 962–973
- [36] Huang, W. et al. (2023). Goal-guided transformer-enabled reinforcement learning for efficient autonomous navigation. *arXiv*. Published online January 1, 2023. <http://arxiv.org/abs/2301.00362>
- [37] Vallacher, R.R. and Wegner, D.M. (1987) What do people think they’re doing? action identification and human behavior. *Psychol. Rev.* 94, 3

[38] Rmus, M. et al. (2023) Choice type impacts human reinforcement learning. *J. Cognit. Neurosci.* 35, 314–330

[39] Hornsby, A.N. et al. (2020) Conceptual organization is revealed by consumer activity patterns. *Comput. Brain Behav.* 3, 162–173

[40] Hommel, B. et al. (2001) The theory of event coding (TEC): A framework for perception and action planning. *Behav. Brain Sci.* 24, 849–878

[41] Fogassi, L. et al. (2005) Parietal lobe: From action organization to intention understanding. *Science* 308, 662–667

[42] Aberbach-Goodman, S. et al. (2022) Same action, different meaning: neural substrates of action semantic meaning. *Cereb. Cortex* 32, 4293–4303

[43] Hommel, B. (1993) Inverting the simon effect by intention: Determinants of direction and extent of effects of irrelevant spatial information. *Psychol. Res.* 55, 270–279

[44] Cheng, S. et al. (2023) Intention beyond desire: Spontaneous intentional commitment regulates conflicting desires. *Cognition* 238, 105513

[45] Chu, J. and Schulz, L.E. (2020) Play, curiosity, and cognition. *Ann. Rev. Dev. Psychol.* 2, 317–343

[46] Hull, C.L. (1934) The rat's speed-of-locomotion gradient in the approach to food. *J. Comp. Psychol.* 17, 393–422

[47] Reppert, T.R. et al. (2015) Modulation of saccade vigor during value-based decision making. *J. Neurosci.* 35, 15369–15378

[48] Summerside, E.M. et al. (2018) Vigor of reaching movements: reward discounts the cost of effort. *J. Neurophysiol.* 119, 2347–2357

[49] Shadmehr, R. and Ahmed, A.A. (2020) *Vigor: Neuroeconomics of Movement Control*. MIT Press

[50] Eisenberger, R. et al. (1976) Learned industriousness and social reinforcement. *J. Pers. Soc. Psychol.* 33, 227–232

[51] Shenhav, A. et al. (2013) The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* 79, 217–240

[52] Colas, C. et al. (2022) Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: A short survey. *J. Artif. Intell. Res.* 74, 1159–1199

[53] Keramati, M. and Gutkin, B. (2014) Homeostatic reinforcement learning for integrating reward collection and physiological stability. *eLife* 3, e04811

[54] Siep, N. et al. (2009) Hunger is the best spice: An fMRI study of the effects of attention, hunger and calorie content on food reward processing in the amygdala and orbitofrontal cortex. *Behav. Brain Res.* 198, 149–158

[55] Levy, D.J. et al. (2013) State dependent valuation: The effect of deprivation on risk preferences. *PLoS One* 8, e53978

[56] Minamimoto, T. et al. (2009) Measuring and modeling the interaction among reward size, delay to reward, and satiation level on motivation in monkeys. *J. Neurophysiol.* 101, 437–447

[57] Juechems, K. et al. (2019) A network for computing value equilibrium in the human medial prefrontal cortex. *Neuron* 101, 977–987.e3

[58] O'Reilly, R.C. (2020) Unraveling the mysteries of motivation. *Trends Cogn. Sci.* 24, 425–434

[59] Bavar, S. et al. (2018) Reference-point centering and range-adaptation enhance human reinforcement learning at the cost of irrational preferences. *Nat. Commun.* 9, 4503

[60] Bavar, S. et al. (2021) Two sides of the same coin: Beneficial and detrimental consequences of range adaptation in human reinforcement learning. *Sci. Adv.* 7, eabe0340

[61] Molinaro, G. and Collins, A.G.E. (2023) Intrinsic rewards explain context-sensitive valuation in reinforcement learning. *PLoS Biol.* 21, e3002201

[62] Markle, A. et al. (2018) Goals as reference points in marathon running: A novel test of reference dependence. *J. Risk Uncertainty* 56, 19–50

[63] Heath, C. et al. (1999) Goals as reference points. *Cognit. Psychol.* 38, 79–109

[64] Dayan, P. and Berridge, K.C. (2014) Model-based and model-free pavlovian reward learning: revaluation, revision, and revelation. *Cogn. Affect. Behav. Neurosci.* 14, 473–492

[65] Robinson, M.J. and Berridge, K.C. (2013) Instant transformation of learned repulsion into motivational “wanting”. *Curr. Biol.* 23, 282–289

[66] Melnikoff, D.E. and Bailey, A.H. (2018) Preferences for moral vs. immoral traits in others are conditional. *Proc. Natl. Acad. Sci.* 115, E592–E600

[67] Frömer, R. et al. (2019) Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nat. Commun.* 10, 4926

[68] Sepulveda, P. et al. (2020) Visual attention modulates the integration of goal-relevant evidence and not value. *eLife* 9, e60705

[69] Harlow, H.F. (1950) Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys. *J. Comp. Physiol. Psychol.* 43, 289–294

[70] Bromberg-Martin, E.S. and Hikosaka, O. (2009) Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63, 119–126

[71] Charpentier, C.J. et al. (2018) Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci.*

[72] Chentanez, N. et al. (2004) Intrinsically motivated reinforcement learning. In *Adv. Neural Inf. Process. Syst.*, vol. 17. MIT Press

[73] Singh, S. et al. (2010) Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Trans. Auton. Mental Develop.* 2, 70–82

[74] Blain, B. and Sharot, T. (2021) Intrinsic reward: potential cognitive and neural mechanisms. *Curr. Opin. Behav. Sci.* 39, 113–118

[75] Barron, H.C. et al. (2013) Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat. Neurosci.* 16, 1492–1498

[76] O'Doherty, J.P. et al. (2021) The hierarchical construction of value. *Curr. Opin. Behav. Sci.* 41, 71–77

[77] McDougle, S.D. et al. (2022) Executive function assigns value to novel goal-congruent outcomes. *Cereb. Cortex* 32, 231–247

[78] Cushman, F. and Morris, A. (2015) Habitual control of goal selection in humans. *Proc. Natl. Acad. Sci.* 112, 13817–13822

[79] Fine, J.M. and Hayden, B.Y. (2022) The whole prefrontal cortex is premotor cortex. *Philos. Trans. R. Soc. B* 377, 20200524

[80] Gollwitzer, P.M. (1986) Action phases and mind-sets. In *Handbook of motivation and cognition: Foundations of social behavior* (Sorrentino, R.M. and Higgins, E.T., eds.), pp. 53–92. The Guilford Press

[81] Jara-Ettinger, J. et al. (2016) The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends Cogn. Sci.* 20, 589–604

[82] Moskowitz, G.B. and Grant, H., eds. (2009) *The psychology of goals*. Guilford Press

[83] Pezzulo, G. et al. (2018) Hierarchical active inference: A theory of motivated control. *Trends Cogn. Sci.* 22, 294–306

[84] Davidson, G. et al. (2022) Creativity, compositionality, and common sense in human goal generation. *Proc. Annu. Meet. Cogn. Sci. Soc.* 44

[85] Baranes, A.F. et al. (2014) The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Front. Neurosci.* 8

[86] Ten, A. et al. (2021) Humans monitor learning progress in curiosity-driven exploration. *Nat. Commun.* 12, 5972

[87] Poli, F. et al. (2022) Contributions of expected learning progress and perceptual novelty to curiosity-driven exploration. *Cognition* 225, 105119

[88] Kool, W. and Botvinick, M. (2018) Mental labour. *Nat. Hum. Behav.* 2, 899–908

[89] Sakaki, M. et al. (2023) Motivated for near impossibility: How task type and reward modulate task enjoyment and the striatal activation for extremely difficult task. *Cogn. Affect. Behav. Neurosci.* 23, 30–41

[90] Son, L.K. and Metcalfe, J. (2000) Metacognitive and control strategies in study-time allocation. *J. Exp. Psychol. Learn. Mem. Cogn.* 26, 204–221

[91] Wang, J.J. and Bonawitz, E. (2022) Children’s sensitivity to difficulty and reward probability when deciding to take on a task. *J. Cogn. Dev.* 0, 1–13

[92] Botvinick, M.M. (2012) Hierarchical reinforcement learning and decision making. *Curr. Opin. Neurobiol.* 22, 956–962

[93] Ribas-Fernandes, J.J.F. et al. (2011) A neural signature of hierarchical reinforcement learning. *Neuron* 71, 370–379

[94] Biderman, N. and Shohamy, D. (2021) Memory and decision making interact to shape the value of unchosen options. *Nat. Commun.* 12, 4648

[95] Aarts, H. et al. (2004) Goal contagion: perceiving is for pursuing. *J. Pers. Soc. Psychol.* 87, 23–37

[96] Doeberl, S. and Munakata, Y. (2018) Group influences on engaging self-control: Children delay gratification and value it more when their in-group delays and their out-group doesn’t. *Psychol. Sci.* 29, 738–748

[97] King, R.B. and Mendoza, N.B. (2020) Achievement goal contagion: mastery and performance goals spread among classmates. *Soc. Psychol. Educ.* 23, 795–814

[98] Baker, C.L. et al. (2009) Action understanding as inverse planning. *Cognition* 113, 329–349

[99] Whiten, A. et al. (2006) Imitation of hierarchical action structure by young children. *Dev. Sci.* 9, 574–582

[100] Call, J. et al. (2005) Copying results and copying actions in the process of social learning: chimpanzees (*pan troglodytes*) and human children (*homo sapiens*). *Anim. Cogn.* 8, 151–163

[101] Charpentier, C.J. et al. (2020) A neuro-computational account of arbitration between choice imitation and goal emulation during human observational learning. *Neuron* 106, 687–699.e7

[102] Losin, E.A.R. et al. (2015) Brain and psychological mediators of imitation: sociocultural versus physical traits. *Cult. Brain* 3, 93–111

[103] Matz, D.C. and Hinsz, V.B. (2000) Social comparison in the setting of goals for own and others' performance. *J. Bus. Psychol.* 14, 563–572

[104] Willis, G.B. and Guinote, A. (2011) The effects of social power on goal content and goal striving: A situated perspective. *Soc. Personal. Psychol. Compass* 5, 706–719

[105] Taylor-Davies, M. et al. (2023) Selective imitation on the basis of reward function similarity. *CogSci* 45

[106] Moty, K. and Rhodes, M. (2021) The unintended consequences of the things we say: What generic statements communicate to children about unmentioned categories. *Psychol. Sci.* 32, 189–203

[107] Ng, A.Y. and Russell, S.J. (2000) Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML, pp. 663–670. Morgan Kaufmann Publishers Inc.

[108] Hadfield-Menell, D. et al. (2016). Cooperative inverse reinforcement learning. *arXiv*. Published online November 12, 2016. <https://doi.org/10.48550/arXiv.1606.03137>

[109] Christian, B. (2020) *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton & Company

[110] Csikszentmihalyi, M. (2009) *Flow: The psychology of optimal experience*. Harper Perennial Modern Classics. Harper and Row

[111] Berlyne, D.E. (1960) *Conflict, arousal, and curiosity*. Conflict, arousal, and curiosity. McGraw-Hill Book Company

[112] Kidd, C. et al. (2012) The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One* 7, e36399

[113] Schmidhuber, J. (1991) A possibility for implementing curiosity and boredom in model-building neural controllers. In *From Animals to Animats* (Meyer, J.A. and Wilson, S.W., eds.). The MIT Press

[114] Chew, B. et al. (2021) A neurocomputational model for intrinsic reward. *J. Neurosci.* 41, 8963–8971

[115] Ryan, n. and Deci, n. (2000) Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemp. Educ. Psychol.* 25, 54–67

[116] Colas, C. et al. (2019). CURIOUS: Intrinsically motivated modular multi-goal reinforcement learning. *arXiv*. Published online May 29, 2019. <http://arxiv.org/abs/1810.06284>

[117] Melnikoff, D.E. et al. (2022) A computational theory of the subjective experience of flow. *Nat. Commun.* 13, 2252

[118] Du, Y. et al. (2022) The relationship between habits and motor skills in humans. *Trends Cogn. Sci.* 26, 371–387

[119] Suits, B. (2014) *The Grasshopper: Games, Life and Utopia*. Broadview Press, 3rd ed.

[120] Gopnik, A. (2020) Childhood as a solution to explore–exploit tensions. *Philos. Trans. R. Soc. B* 375, 20190502

[121] Dobzhansky, T. (1973) Nothing in biology makes sense except in the light of evolution. *Am. Biol. Teach.* 35, 125–129

[122] Sumers, T.R. et al. (2023) Show or tell? exploring when (and why) teaching with language outperforms demonstration. *Cognition* 232, 105326

[123] Colas, C. et al. (2020) Language as a cognitive tool to imagine goals in curiosity driven exploration. *Adv. Neural Inf. Process. Syst.* 33, 3761–3774

[124] Schultz, W. (2007) Reward. *Scholarpedia* 2, 1652

[125] Schaul, T. et al. (2015) Universal value function approximators. In *ICML*, pp. 1312–1320

[126] Duncan, J. (2013) The structure of cognition: Attentional episodes in mind and brain. *Neuron* 80, 35–50

[127] Wulf, G. and Lewthwaite, R. (2016) Optimizing performance through intrinsic motivation and attention for learning: The OPTIMAL theory of motor learning. *Psychon. Bull. Rev.* 23, 1382–1414

[128] Chauvel, G. et al. (2015) Visual illusions can facilitate sport skill learning. *Psychon. Bull. Rev.* 22, 717–721

[129] Katzir, M. et al. (2020) Cognitive performance is enhanced if one knows when the task will end. *Cognition* 197, 104189

[130] Saemi, E. et al. (2012) Knowledge of results after relatively good trials enhances self-efficacy and motor learning. *Psychol. Sport Exerc.* 13, 378–382

[131] Sayali, C. et al. (2023) Learning progress mediates the link between cognitive effort and task engagement. *Cognition* 236, 105418

[132] Chu, J. and Schulz, L. (2022) “Because I want to”: Valuing goals for their own sake. *Proc. Annu. Meet. Cogn. Sci. Soc.* 44

[133] Ho, M.K. et al. (2023). Rational simplification and rigidity in human planning. *PsyArXiv*. Published online March 30, 2023. <https://doi.org/10.31234/osf.io/aqxws>

[134] Kruglanski, A.W. et al. (2018) A structural model of intrinsic motivation: On the psychology of means-ends fusion. *Psychol. Rev.* 125, 165–182

[135] Schwartz, B. (2015) *Why We Work*. Simon and Schuster

[136] Santucci, V.G. et al. (2016) GRAIL: A goal-discovering robotic architecture for intrinsically-motivated learning. *IEEE Trans. Cogn. Develop. Syst.* 8, 214–231

[137] Klinger, E. (2013) Goal commitments and the content of thoughts and dreams: basic principles. *Front. Psychol.* 4

[138] Ballard, I. et al. (2023). Reward reinforcement creates habitual selection of goals. *PsyArXiv*. Published online February 23, 2023. <https://doi.org/10.31219/osf.io/63aqz>

- [139] Dezfouli, A. and Balleine, B.W. (2013) Actions, action sequences and habits: Evidence that goal-directed and habitual action control are hierarchically organized. *PLOS Comput. Biol.* 9, e1003364
- [140] Keramati, M. et al. (2017) Cocaine addiction as a homeostatic reinforcement learning disorder. *Psychol. Rev.* 124, 130–153