CellPress REVIEWS

Review

Intuitive Physics: Current Research and Controversies

James R. Kubricht,[1,*] Keith J. Holyoak,[1] and Hongjing Lu[1]

Early research in the field of intuitive physics provided extensive evidence that humans succumb to common misconceptions and biases when predicting, judging, and explaining activity in the physical world. Recent work has demonstrated that, across a diverse range of situations, some biases can be explained by the application of normative physical principles to noisy perceptual inputs. However, it remains unclear how knowledge of physical principles is learned, represented, and applied to novel situations. In this review we discuss theoretical advances from heuristic models to knowledge-based, probabilistic simulation models, as well as recent deep-learning models. We also consider how recent work may be reconciled with earlier findings that favored heuristic models.

Trends

People demonstrate common misconceptions when asked to make explicit reasoning judgments about physical systems on pencil-and-paper tasks. However, their performance improves when the problem is accompanied with dynamic and contextual information.

Recent research has shown that our implicit judgments about physical situations based on rich dynamic displays are consistent with probabilistic mental simulation governed by normative physical principles, but are subject to biases evoked by perceptual uncertainty and prior beliefs about physical variables.

Recent behavioral evidence suggests that people utilize a cognitive intuitive physics engine to reason about a broad range of physical situations.

Researchers have begun to explore how to integrate probabilistic mental simulation with deep-learning models to extract higher-level physical knowledge from static and dynamic visual inputs.

Reinvigorating Intuitive Physics

Humans are able to understand their physical environment and interact with objects and substances that undergo dynamic state changes, making at least approximate predictions about how observed events will unfold (e.g., predicting the trajectory of a thrown ball, the direction that a chopped tree will fall, or the path of a breaking wave). The knowledge underlying such activities is termed **intuitive physics** (see Glossary). This topic, an active research area for several decades, has recently been reinvigorated by new theoretical approaches linked to artificial intelligence. These theories have been used to model findings from behavioral studies that apply psychophysical measures to perception and reasoning with complex dynamic displays. Here, we review recent research and theories (placing them in the context of earlier work), and discuss some of the ongoing controversies in the field.

Apparent Misconceptions in Intuitive Physics

Before the most recent decade, research on intuitive physics primarily focused on misconceptions that people demonstrate when reasoning about the attributes and movements of objects and substances in the world (e.g., [1,2]). Numerous studies found that humans exhibit striking deviations from Newtonian physical principles when asked to explicitly reason about the expected continuation of a dynamic event based on a static image representing the situation at a single timepoint. The predictions that people made in these studies often appeared to agree with erroneous theories of motion, rather than with (ground-truth) **Newtonian physics** (Box 1). For instance, adults often predict that an object dropped from a moving body will follow a linear path downwards (i.e., the straight-down belief [2]), and children predict that a horizontal force will propel a vertically moving object in the direction that it is pushed [3]. Such evidence has been used to argue that people sometimes reason about the physical world using an Aristotelian model of physics. In other situations, adults appear to exhibit medieval impetus beliefs, for example, that an object exiting a curved tube will follow a curvilinear trajectory in the absence of external forces [1]. In these studies (Figure 1 and Box 2), participants were typically shown physical situations at a specific point in time via static diagrams printed on paper, and were required to draw how the situation would unfold going forwards in time. In such

[1]Department of Psychology, University of California, Los Angeles, 405 Hilgard Avenue, Los Angeles, CA 90095, USA

*Correspondence: kubricht@ucla.edu (J.R. Kubricht).

CellPress
REVIEWS

Box 1. Erroneous Theories of Motion

The motions of objects and substances in the world are accurately described by the principles of Newtonian physics. Newton's three laws state that (i) an object at rest will stay at rest, and an object in motion will stay in motion, unless acted upon by an external force, (ii) objects accelerate when an external force is applied to them, and (iii) for every action, there is an equal and opposite reaction. Therefore, when objects fall, a gravitational force accelerates them downwards, and when a force is applied to an object in a given direction, the object retains the motion it had before the force was exerted.

By contrast, medieval impetus theory states that (i) the act of setting an object in motion imparts upon it an 'impetus' force that is used to maintain its motion, and (ii) the impetus force of a moving object gradually dissipates over time [15]. Thus, an object thrown in the air falls down because its vertical impetus dissipates, and an object will continue along a curved trajectory after it is released from a pendulum because of a (dissipating) curvilinear impetus force.

The principles of Aristotelian physics state that an object will move in the direction that it is pushed [3]. In other words, if a horizontally moving object is given a vertical push (e.g., a gravitational force downwards), it will immediately move straight downwards and lose its horizontal component of velocity.

pencil-and-paper tasks, people often succumb to systematic errors when predicting the physical behavior of situations. People are not, however, internally consistent in their explicit reasoning judgments. Instead, they appear to reason in accord with different theories of motion depending on the situation [4]. These findings led to a generally pessimistic assessment of the human capacity to perceive and reason about physical situations, most notably in projectile motion and object collision situations [5,6].

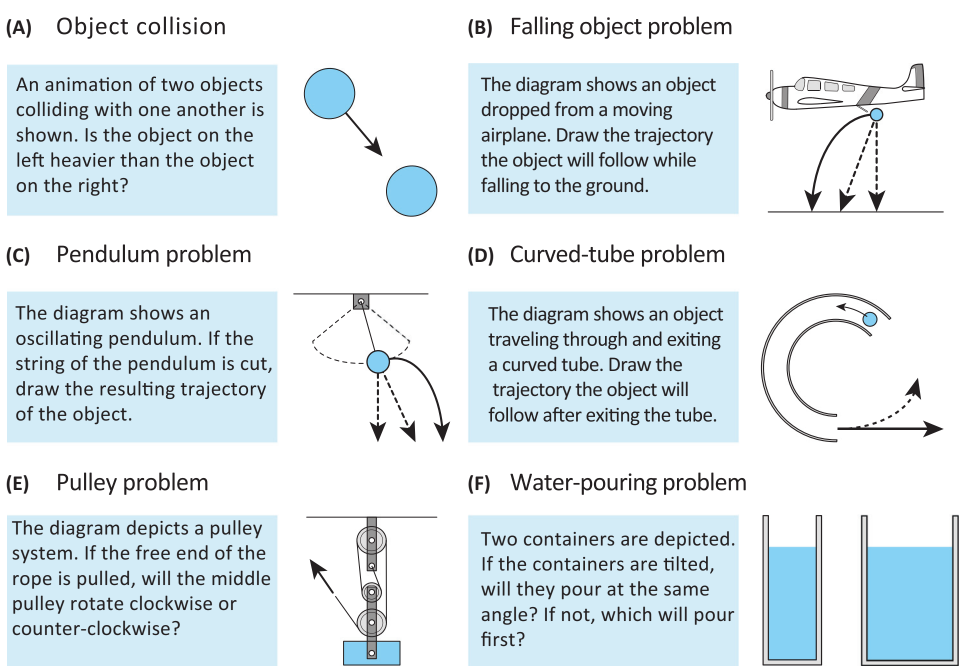

Trends in Cognitive Sciences

Figure 1. Examples of Intuitive Physics Problems. The task in each problem is to reason about the attributes or movements of objects and substances in various situations. Aside from object collision judgments (A), problems have generally been depicted via a static diagram of the physical system. In (B–D) the unbroken line corresponds to the correct trajectory, and the broken lines correspond to common, erroneous predictions. The probabilistic simulation framework has achieved success predicting people's expectations about the attributes (A) and movements (C) of objects in dynamic displays, as well as the pouring angle of two fluid-filled containers (F). A computational account of people's explicit trajectory predictions in (B–D) has not been developed.

Glossary

Convolutional neural network: a neural network model in which lower levels are defined by parameters consisting of a set of learnable filters that are convolved across the width and height of the input image to extract higher-order pattern activations.

Deep learning: a hierarchical machine-learning method composed of multiple processing layers that extract abstract representations from low-level features such as pixel intensity and color in 2D images.

Dynamics: a branch of Newtonian mechanics that studies how forces (linear and rotational) affect the movement of objects.

Intuitive physics: the knowledge underlying the human ability to understand the physical environment and interact with objects and substances that undergo dynamic state changes, making at least approximate predictions about how observed events will unfold.

Mass: an attribute of a physical body that serves as a measure of its resistance to acceleration and determines its gravitational weight.

Newtonian physics: a framework based on Newton's three laws which describe the relationships among objects in the physical world, the forces acting upon them, and the motion resulting from those forces.

Object state: specifies the location, movement, and attributes of an object at a given time.

Box 2. Early Research on Intuitive Physics

Early research in the field provides several examples of situations in which humans demonstrate common misconceptions about how objects in the environment behave. The most extensive line of research involves human judgments about object collisions, specifically about which object appears heavier based on an animated depiction of the collision event. Interestingly, people have a tendency to believe that the object that moves with a greater initial velocity is heavier [5]. This bias has recently been reinterpreted within the probabilistic simulation framework [13].

Another major line of research examined human knowledge about projectile motion. In these tasks, participants viewed a static depiction of an object moving in a given situation and were asked to draw the trajectory the object would follow as time progressed. In the falling object problem, people commonly responded that an object dropped from a moving body will follow a straight path downwards [2]. People's predictions about the trajectory of an object after being released from a pendulum were also inconsistent with Newtonian principles [57], but a probabilistic simulation model has achieved success predicting where the object is expected to land. People also commonly predict that an object exiting a curved tube will follow a nonlinear trajectory upon exiting [58]. Although people's trajectory predictions in these problems are inconsistent with normative physical principles, they correctly judge anomalous trajectories as appearing unnatural when presented in an animated format [9]. Currently, no computational account explains people's erroneous misconceptions in explicit trajectory prediction problems.

Work on mechanical reasoning supports the role of mental simulation in intuitive physics. In the pulley problem, people take longer to respond for pulleys farther from the beginning of the causal sequence (i.e., the rope being pulled [27]). This finding suggests that humans build analog representations of physical situations that carry spatial information, as opposed to reasoning according to a simple rule (e.g., alternating pulleys rotate in opposite directions). The pervasiveness of mental simulation is further reinforced by results from the water-pouring task [28]. In this task, participants were successful when forming predictions using mental simulation but performed much worse when solving the problem using explicit reasoning. A probabilistic simulation model has recently accounted for people's mental simulation performance in a modified water-pouring task [38]. Taken together, these problems provide several examples in which people's explicit reasoning about the physical world is biased, and suggest experimental factors that lead to accurate predictions and judgments.

However, several studies have shown that these misconceptions can be reduced or even abolished when experimental paradigms and tasks are varied in particular ways [7–11]. Human errors are greatly reduced when explicit reasoning problems are presented in a familiar context (e.g., when an object exiting a curved tube is replaced by water exiting a curved hose [7]). This finding suggests that specific prior knowledge can override more general physical intuitions that generate misconceptions. In addition, recent work has demonstrated that, although adults perform poorly when drawing the trajectory an object will follow after being released from a pendulum, they can successfully predict its landing location [8]. It thus appears that systematic misconceptions are less likely to be exhibited in tasks that evoke implicit or tacit knowledge, such as recognizing the normative unfolding of an event in an animated display. In addition, the format of the stimulus display also influences how susceptible people are to misconceptions about object movements. For example, human judgments are more consistent with Newtonian physics when situations are presented in an animated format [9,10]. Developmental studies have also yielded converging evidence that the judgments of children are more Newtonian when events are presented as animations [11].

In addition to studies involving motion of a single object, early research about collisions between two objects also demonstrated deviations of human judgments from what is expected given Newtonian principles (Box 3). These studies were inspired by classic work on perceptual causality which found that people report causal impressions that are not fully determined by the physical properties of the situation [12,13]. For example, consider the case in which an initially moving object (motor object) collides with one that is initially stationary (projectile object). When the physical effect of the motor object on the projectile object is relatively small (e.g., the post-collision velocity of the projectile object is less than the pre-collision velocity of the motor object), people report a stronger causal impression than when the physical effect is large (e.g., the post-collision velocity of the projectile object is greater than the pre-collision velocity of the motor object) [13]. In addition, when two equally heavy objects collide, people often report that the

Box 3. Object Collision Physics

In the case of collision events, two objects move towards one another, collide, and then move away from one another. The movement of each object both before and after the collision is represented by its velocity, which specifies the speed and direction of each object. While velocity can be perceived, the mass of each object (i.e., how heavy each object is) and the amount of restitution (i.e., how much energy is 'given back' in the collision) cannot.

The momentum of an object is defined by the product of its mass and velocity. The principle of conservation of momentum states that, in a closed system, the sum of the momentum of objects before a collision is equal to the sum of their momentum after (i.e., momentum is conserved). By rearranging the conservation of momentum expression for two-body collisions, the relative mass of two objects can be expressed purely in terms of the initial and final velocities of the objects [59]. Thus, if humans can perceive the velocity of objects without error and reason in accord with momentum conservation, their relative mass judgments should be invariant to the amount of restitution in the collision. This is termed the direct perception model. Under this view, people should be equally accurate when reasoning about the relative masses of two billiard balls (high restitution) as they are when reasoning about those of two rubber balls (low restitution). However, humans are increasingly inaccurate when reasoning about relative mass in low-restitution collisions, thus deviating from what would be expected given Newtonian physical principles.

motor object weighs more than the projectile object, a phenomenon termed the motor object bias [5]. This finding was traditionally interpreted in terms of heuristics: people may infer the attributes of colliding objects using two rules based on salient perceptual cues. (i) The object that moves fastest following a collision event is lighter (the velocity heuristic; Figure 2A, Key Figure); and (ii) the object that deflects at the greatest angle is lighter (the angle heuristic) [6].

However, although these heuristics account for human judgments about the relative **mass** of colliding objects in some cases, they do not generalize to other situations. For example, it is unclear whether the heuristic model accounts for relative mass judgments when the motion of each object before the collision is occluded [13]. The heuristic account also fails to explain why people are less susceptible to the motor object bias after completing a large number of training trials. Even if naïve observers use heuristic reasoning to make perceptual judgments, with experience they may transition to correct application of normative physical principles [14].

The general picture that emerges from research on intuitive physics is that people exhibit misconceptions and biases when (i) they are asked to provide explicit predictions or explanations about continuations of physical events, (ii) the events are unfamiliar and presented with minimal context, and (iii) the events are portrayed using impoverished stimuli, such as static line drawings depicting a situation at a single moment in time. Although misconceptions have often been attributed to people holding erroneous Aristotelian or impetus theories, it may well be the case that under these unfavorable conditions people do not employ a systematic theory to infer physical motion [4]. Instead, explicit predictions may draw upon a set of individualized background knowledge [8] based on salient perceptual cues that seem to be potentially relevant. Importantly, when these unfavorable conditions are alleviated (i.e., when people make more tacit judgments about familiar types of events based on rich visual displays), human judgments align more closely with Newtonian physics [7].

It thus seems possible that people in fact have a strong intuitive 'physics engine' available to them, but it is only evoked under favorable conditions. This does not imply, however, that people do not hold explicit conceptions of physical processes, nor that those conceptions cannot interfere with human reasoning in familiar environments. Indeed, people sometimes push objects along curved paths to impart upon them a curvilinear impetus [15], and even erroneously describe their own actions as adhering to the straight-down belief [2]. However, cortical activation associated with explicit physical knowledge [16] does not entirely overlap with brain activities associated with tacit physical inference [17]. Although people might plan their movements or describe their actions in accord with explicit physical conceptions, it

Key Figure

Three Approaches to Relative Mass Judgment

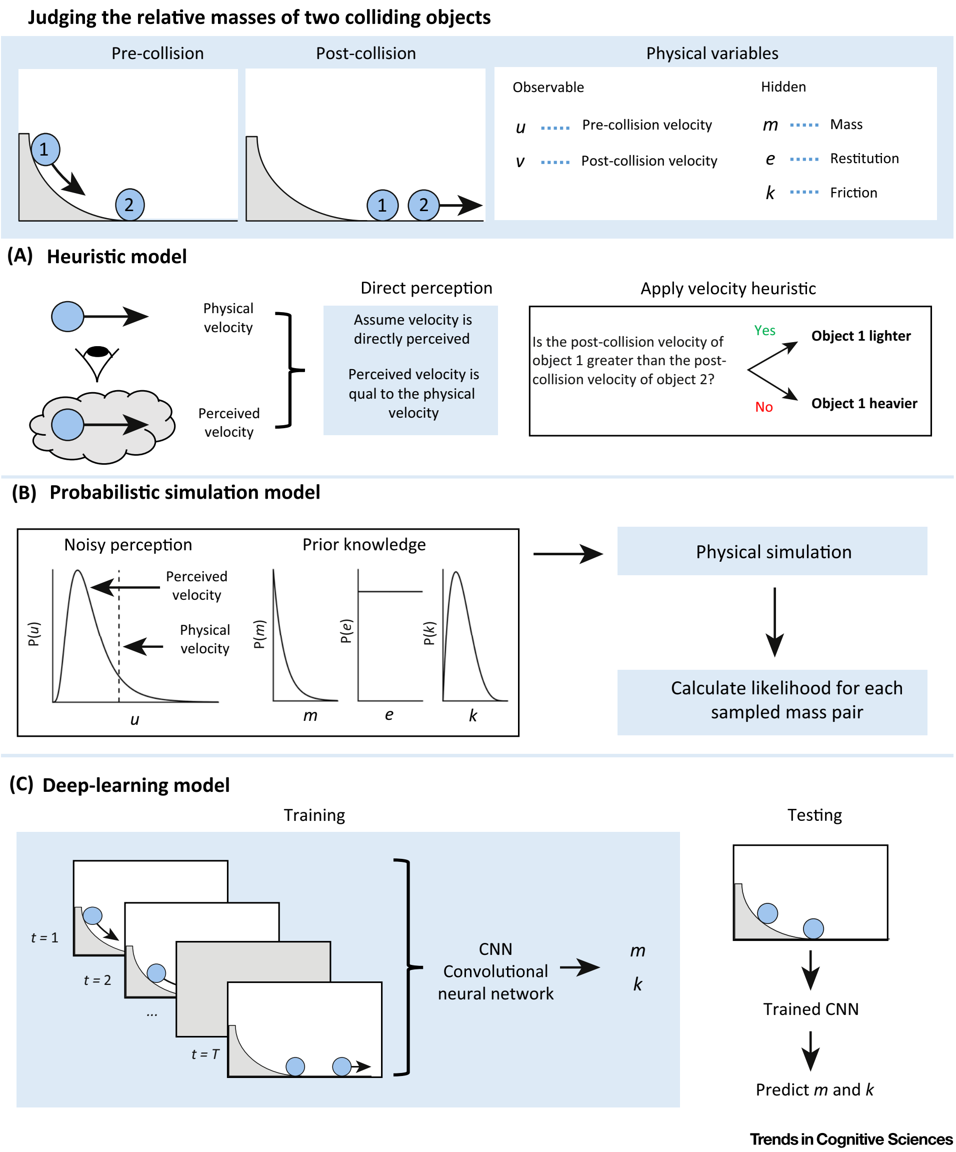

Figure 2. Description of three computational approaches to determining the relative masses of two colliding objects. The main differences between the models involve the role of learning (minimal for the heuristic approach, limited for probabilistic simulation, and central for deep learning). (A) In the heuristic model, it is assumed that observed velocities are equivalent to physical velocities in the world (i.e., velocity is directly perceived). Post-collision velocities are compared, and the object that moves at the greatest speed following collision is assumed to be lighter. This heuristic is not attributed to learning. (B) The probabilistic simulation model places priors on hidden physical variables. The motion prior biases perceived velocities towards slow motion. The likelihoods of different mass ratios are determined by comparing simulated final velocities to observed velocities. Learning may impact on the prior knowledge involved in the inference. (C) In a deep-learning model, a convolutional neural network (CNN) is trained on 2D image inputs and outputs object attributes (mass and friction). The CNN is then used to predict object attributes from previously unseen image data. This approach uses bottom-up learning.

appears that their tacit judgments and predictions are generally consistent with normative physical principles [8,10].

It remains uncertain how explicit physical conceptions are derived from experience and to what degree they interact with tacit physical knowledge [18]. One reason for this uncertainty is the difficulty in classifying conceptions as arising from perceptual ambiguity or from ineffective representations of the perceptual and physical variables involved in a task. For example, the straight-down belief may appear to arise when drawing the trajectory of an object after being released from a pendulum because the speed of the object at the indicated (static) location is ambiguous. Alternatively, people might believe that an object dropped from a moving body will fall straight downwards because they represent the perceived motion of the object relative to the moving body. The finding that humans are less susceptible to the straight-down belief when the moving body is removed from the situation [19] is consistent with the latter possibility. People are also commonly biased when drawing the water level on a rotated container (i.e., the water-level problem [20]), even after explicitly stating that the surface of a liquid should remain horizontal regardless of the orientation of its container [21]. In this case, the bias appears to arise due to ineffective representation: specifically using an object-centered reference frame with axes parallel to the surfaces of the container [22]. Although such biases might be interpreted as predictions based on erroneous physical theories, they might instead reflect correct application of normative principles to variables represented in moving or rotated frames of reference. It is also possible that meaningful physical quantities (e.g., the mass distribution of a rotating wheel) are simply not represented or are coarsely approximated in some situations, leading to erroneous predictions and judgments [23]. People's judgments about physical quantities (e.g., the forces that two colliding objects impart upon one another) have also been shown to disagree with fundamental Newtonian principles [24]. Thus, it is important that future work in intuitive physics should consider (i) correspondence between cognitive and physical constructs, (ii) the nature of cognitive representations across dissimilar problem contexts, (iii) the role of physical approximation in complex displays, and (iv) the interaction between explicit conceptions and tacit understanding in prediction and judgment tasks.

Noisy Newton Framework

In recent years, research on intuitive physics has been reinvigorated by new theoretical approaches based on Bayesian inference, most notably the noisy Newton framework, which integrates ground-truth physical principles with uncertainty about sensory information [13]. Models based on the noisy Newton framework assume that people integrate noisy sensory inputs with prior beliefs about perceptual and physical variables underlying physical situations, and model the constraints among those variables in accord with Newtonian physics. In the case of collision events, predictions are modeled by simulating thousands of physical situations. In each simulation, the physical outcome is computed using Newtonian laws operating on sampled variables of perceptual and physical properties. Although most perceptual variables appear to be observable (e.g., velocity, location), it remains necessary to convert objective evidence (observations) into subjective estimates by integrating noisy sensory input with priors on statistical regularities of perceptual cues in the world. Another source of sampling uncertainty comes from unobservable physical properties (e.g., mass, viscosity) which must be inferred from sensory observations and/or general knowledge about the physical world (Figure 2B). The noisy Newton framework effectively reconciles several inconsistencies between human judgments and Newtonian physics [13].

Probabilistic Simulation Approaches to Intuitive Physics

Motivated by the initial successes of the framework, several researchers have recently extended the noisy Newton approach to explain human judgments and predictions about a variety of physical situations. As an overarching framework, the Bayesian approach to intuitive physics can be viewed as a tool for understanding how abstract knowledge – represented by priors and a generative function to assess likelihood – guides inference about **object states** from incomplete and noisy information about perceptual and physical variables [25]. In the noisy Newton framework, inference is achieved by passing noisy information through a physics engine, which is defined by the principle of conservation of momentum in the case of object collisions [13]. Knowledge about object **dynamics** is 'written into' the model under the assumption that the transformation of perceptual inputs into physical expectations is in accord with Newtonian physical constraints.

The key idea underlying probabilistic simulation models is that humans construct mental models about physical situations, allowing for inference of future object states through mental simulation [26]. The role of mental simulation is supported by work on mechanical reasoning which has demonstrated that people reason about physical systems by constructing and transforming spatial representations to answer questions about the behavior of objects and substances [27,28]. Spatial representation implies that object locations, motions, and hidden attributes in the physical world – as well as their interactions – are encoded and represented in the mind [29]. Recent neural evidence suggests that the mental simulation process is likely carried out in cortical regions that overlap with the domain-general 'multiple demand' system of the brain [17].

Probabilistic simulation models make judgments in physical reasoning tasks by integrating noisy information processing with advanced physics-based graphics engines to simulate future object states. In each simulation, the values of perceptual and physical variables in a scene are sampled according to distributions that emulate noisy information processing of the positions, velocities, and attributes of the objects. Based on sampled states of perceptual and physical inputs, an 'intuitive physics engine' which approximates Newtonian principles is used to simulate future object states. The outcome of each simulation is then queried to form a predicted judgment, such as whether or not a tower of blocks fell down [26] or how much liquid fell into a designated area [30]. Judgments are then aggregated across simulations to form a predicted response distribution. Parameters in the simulation model are chosen such that the distribution accurately reflects human behavior.

The probabilistic simulation approach has demonstrated promising results across several physical domains. Most studies utilizing probabilistic simulation examine the correlation between human performance and model predictions across a range of experimental conditions, rather than absolute performance levels (the typical focus in earlier studies). Overall, model predictions correlate well with human responses about the motions and attributes of stacked blocks [26,31] and about how liquids move past obstacles [30]. Moreover, it has been shown that causal judgments about object collision outcomes are correlated with counterfactual assessments, specifically whether and how an outcome occurs [32]. The approach has also been used to explain human inferences about whether different containers can hold particular objects [33], as well as infant reasoning about complex displays of moving objects [34]. It is important to note that these simulation models account for human predictions primarily through the implementation of noisy information processing – the physical principles underlying object and substance dynamics are approximated but are not systematically biased.

However, uncertainty in the inferential mechanism (i.e., the physical model itself) may also constitute an important component in the intuitive physics 'module' of the brain. It has been shown that implementing randomness into the dynamics of objects in physical situations (on top of perceptual noise) provides a better fit to human performance in predicting the trajectory of an occluded object bouncing within a box [35]. It appears that mental simulation outcomes are not deterministic for humans; instead, intermediate object states are randomly perturbed during the inference process. Recent work suggests that, as the outcome of an event becomes increasingly uncertain or a problem becomes more difficult, additional cognitive resources are allocated to the mental simulation process [36].

Intuitive Physics with Liquids and Other Substances

Research on intuitive physics has transitioned beyond the behavior of solid objects to examine the human capacity to reason about the dynamics of liquids. This ability appears to emerge in the earliest stages of life because 5-month-old infants are able to distinguish between solids and non-solid substances [37]. A recent study showed that a computational model based on probabilistic simulation can account for human performance in predicting the resting configuration of a liquid pouring past obstacles into two empty basins [30]. This model included perceptual noise applied to the initial locations of liquid elements, and simulated their movements using approximated normative physical principles. Probabilistic simulation was also employed to model performance in an explicit reasoning task concerning the angle at which a liquid-filled container will begin to spill through mental simulation [38]. This paradigm is a variant of an earlier experiment in which participants reasoned about the angle at which two water-filled containers – one wider than the other – would begin to spill [28] (Figure 1F and Box 2). When asked explicitly to reason about which container would pour at a lesser angle, most participants mistakenly chose to rotate the wider container further. However, when asked to reason about the pouring angle of each container independently through imagined action, participants correctly rotated the narrower container further. In the modified task [38], two containers were filled with fluids varying in their viscosity (i.e., the apparent thickness or stickiness of a fluid), and visual motion cues based on flow visualization animations [39] were used to infer the viscosity of each fluid. Viscosity inferences were consistently biased towards lower values of viscosity, suggesting that people have a prior belief that fluids tend to behave like water. Results from the viscous fluid-pouring task indicated that human judgments about the pouring angle of the two containers are consistent with a probabilistic simulation model utilizing normative physical principles given noisy perceptual inputs.

Inference of Physical Variables

The probabilistic simulation approach builds on two basic components: physical variables provided as the input to a physics engine, and physical principles encoded in the engine. Some physical variables (e.g., velocity and object positions) can be directly perceived, although perceived values could be distorted by neural noise and by generic priors (e.g., the slow and smooth prior in motion perception [40,41]). However, other physical variables (e.g., mass, viscosity, density, and gravity) are not directly perceivable. How could humans infer these physical attributes from low-level visual features of images?

Recent advances in **deep learning** models suggest a potential computational mechanism for inferring physical attributes from visual inputs and making predictions about physical situations. This approach arose in the field of machine learning, and is based on implementations of **convolutional neural networks** (CNNs) [42–44]. These networks take images encoded at the pixel level as inputs, and process the information through hierarchical layers to learn representations at multiple levels of abstraction, ranging from simple visual components (e.g., edges) to more complex patterns and object categories. A hybrid approach – integrating a knowledge-based physics model with a learning-based recognition network for predicting physical attributes from visual inputs – has had some success in accounting for human intuitive physical predictions [45]. Utilizing a deep-learning network, dynamic visual inputs (sequences of 2D images) are mapped to inferred attributes (mass and friction) of two colliding objects through multiple processing layers (Figure 2C). This procedure effectively inverts a key component of the generative physical process. The network is trained on image data tied to object attributes, which are determined by matching key features of the visual inputs to simulation output from a physics engine. This model performs with an accuracy comparable to that of humans, demonstrating that learning-based methods can be effectively integrated with a knowledge-based physics engine to infer the attributes and dynamics of objects in the environment.

Learning about the Physical World

Approaches to physical reasoning based on probabilistic simulation typically assume that ground-truth physical principles are provided as prior knowledge. Thus, it is important to examine how such knowledge can be acquired. Is the mind of a child akin to a blank notebook with minimal information, as Alan Turing surmised [46], or is a cognitive architecture in place that guides the developing mind in learning about the physical properties of objects and the nature of their interactions?

Converging research on infants, children, and non-human primates has provided support for the core knowledge thesis [47], according to which humans utilize separable systems of innate core knowledge, and these serve as building blocks for later learning. The thesis states that core physical principles guide the construction of tacit theories of motion [48]. Perceptual information serves as evidence for preexisting theories, and is represented with increasing complexity as theoretical understanding develops [49]. For example, infants first appear to grasp qualitatively whether a box and table are in contact with one another, and later become sensitive to the quantitative proportion of contact between the two surfaces. This variable identification process is repeated for different phenomena at various developmental stages, leading to piecemeal knowledge about physical situations which does not necessarily transfer to new situations that seem to be unrelated at the surface level [48]. In other words, initial knowledge about the physical world is specific to learned domains [50].

From a computational perspective, one basic approach to learning is based on exemplars. Observed instances of a physical situation are represented as vectors in an *N*-dimensional space linked to corresponding attributes (e.g., whether the motor or projectile object in a collision is heavier [51]). Expected attributes of newly observed instances are predicted by summing similarity measures across instances belonging to each possible classification. However, although the exemplar approach makes sensible predictions within constrained physical regimes by imitating physical knowledge, it fails to generalize to previously unseen regions in the stimulus space.

A more recent learning-based approach has been to directly emulate physical principles using deep-learning methods. The NeuroAnimator model [52] is a neural network that emulates the mapping function which propagates a physical state forwards in time by viewing several instances of physical state transformations. Whereas the probabilistic simulation approach utilizes closed-form physical expressions to propagate scene states forwards in time, the NeuroAnimator model can achieve comparable performance by learning state transition patterns and applying them to previously unseen situations. When a general-purpose engine that emulates the laws of physics was trained on several physical domains, it then generalized to novel systems with different object quantities and relational rules [53]. This is a promising step towards developing a learnable physics engine that can generalize to novel situations in a manner consistent with human abilities. Similarly, the PhysNet model, with an architecture based on a CNN, has achieved success in making physical predictions about simplistic block tower scenarios [54]. After training on artificial scenarios, PhysNet is capable of reasoning about the outcome and future trajectories of both artificial and real-world block tower configurations. It can generalize to new block-tower configurations with a different number of blocks than in the training cases, and yields predictions that are reliably correlated with human responses.

However, the PhysNet model would have difficulty generalizing to situations that are more dissimilar to the training examples (e.g., the trajectory of a thrown object). In addition, unlike developing infants, the model requires several thousands of training examples to abstract basic physical knowledge about the environment. Future research will determine whether learning-based pattern recognition networks can be utilized to extract generalizable physical knowledge from perceptual inputs.

Concluding Remarks and Future Prospects

The field of intuitive physics in the past three decades has benefited from advances on several fronts: stimulus displays (from static diagrams to vivid dynamic animations controlled by computer graphics), computational theory (from heuristic accounts to a parsimonious framework based on probabilistic mental stimulation), and choice of physical situations to study (from a near-exclusive focus on the movement of solid objects to the behaviors of non-rigid fluids). The field now provides a model domain for quantitatively exploring the complex interrelationships between perception and reasoning.

Work guided by knowledge-based and learning-based approaches to intuitive physics suggests that a human intuitive-physics module integrates perceptual and reasoning processes to infer the behavior of physical situations. However, a great deal of future work will be necessary to develop a learnable and generalizable model of intuitive physical inference (see Outstanding Questions). Research on probabilistic simulation indicates that human predictions about physical situations are consistent with probabilistic inference. Nevertheless, such models require a vast number of simulations based on 'hard-coded' normative physical constraints to generate predicted response distributions [8,13,26,30–32]. On the face of it, such computational complexity appears prohibitively demanding for the cognitive system [55]. Moreover, these models do not provide a full account of how physical principles can be learned through experience [56], nor of how such principles might be implemented in neural circuitry [17]. Nonetheless, recent advances constitute progress towards developing a machine that can perceive, reason about, and interact with physical entities.

Outstanding Questions

How proficient are people at reasoning about the dynamics of relatively unfamiliar non-rigid substances (e.g., sand, honey), and what prior knowledge do people adopt about the attributes of unfamiliar substances?

What perceptual characteristics of intuitive physics problems are necessary to enable spatial representation of physical variables and subsequent mental simulation?

What is the role of dynamic uncertainty (i.e., uncertainty in people's internal model of physical dynamics) in mental simulation? What factors cause changes in the amount of uncertainty, and how can stochastic noise at different levels of processing be implemented into numerical physics models?

How do humans acquire the ability to perform mental simulation, and selectively control the use of mental simulation in different situations? What computational constraints and processing constraints are adopted when reasoning about physical situations?

How are intuitive mental models of Newtonian physics implemented in the neural circuitry of the brain?

To what extent can learning-based models (such as deep-learning models and other neural networks) emulate physical knowledge and explain intuitive physics?

Acknowledgments

Support for the present work was provided by a National Science Foundation (NSF) Graduate Research Fellowship to J.R. K. and NSF grant BCS-1353391 to H.L.

References

1. McCloskey, M. *et al.* (1980) Curvilinear motion in the absence of external forces: naive beliefs about the motion of objects. *Science* 210, 1139–1141
2. McCloskey, M. *et al.* (1983) Intuitive physics: the straight-down belief and its origin. *J. Exp. Psychol. Learn. Mem. Cogn.* 9, 636–649
3. DiSessa, A.A. (1982) Unlearning Aristotelian physics: a study of knowledge-based learning. *Cogn. Sci.* 6, 37–75
4. Cook, N.J. and Breedin, S.D. (1994) Constructing naive theories of motion on the fly. *Mem. Cogn.* 22, 474–493
5. Todd, J.T. and Warren, W.H., Jr (1982) Visual perception of relative mass in dynamic events. *Perception* 11, 325–335
6. Gilden, D.L. and Proffitt, D.R. (1994) Heuristic judgment of mass ratio in two-body collisions. *Atten. Percept. Psychophys.* 56, 708–720
7. Kaiser, M.K. *et al.* (1986) Intuitive reasoning about abstract and familiar physics problems. *Mem. Cogn.* 14, 308–312
8. Smith, K.A. *et al.* (2013) Consistent physics underlying ballistic motion prediction. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society*, pp. 3426–3431, Cognitive Science Society
9. Kaiser, M.K. *et al.* (1992) Influence of animation on dynamical judgments. *J. Exp. Psychol. Hum. Percept. Perform.* 18, 669–689
10. Kaiser, M.K. *et al.* (1985) Judgments of natural and anomalous trajectories in the presence and absence of motion. *J. Exp. Psychol. Learn. Mem. Cogn.* 11, 795–803
11. Kim, I.K. and Spelke, E.S. (1999) Perception and understanding of effects of gravity and inertia on object motion. *Dev. Sci.* 2, 339–362
12. Michotte, A. (1963) *The Perception of Causality,* Basic Books
13. Sanborn, A.N. *et al.* (2013) Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychol. Rev.* 120, 411–437
14. Runeson, S. *et al.* (2000) Visual perception of dynamic properties: cue heuristics versus direct-perceptual competence. *Psychol. Rev.* 107, 525–555
15. McCloskey, M. (1983) Intuitive physics. *Sci. Am.* 248, 122–130
16. Mason, R.A. and Just, M.A. (2016) Neural representations of physics concepts. *Psychol. Sci.* 27, 904–913
17. Fischer, J. *et al.* (2016) Functional neuroanatomy of intuitive physical inference. *Proc. Natl. Acad. Sci.* 113, 5072–5081
18. Howe, C. *et al.* (2014) Children's conceptions of physical events: explicit and tacit understanding of horizontal motion. *Br. J Dev. Psychol.* 32, 141–162
19. Kaiser, M.K. *et al.* (1985) The development of beliefs about falling objects. *Percept. Psychophys.* 38, 433–539

20. Piaget, J. and Inhelder, B. (1956) *The Child's Conception of Space,* Routledge and Kegan Paul
21. Howard, I. (1978) Recognition and knowledge of the water-level problem. *Perception* 7, 151–160
22. McAfee, E.A. and Proffitt, D.R. (1991) Understanding the surface orentation of liquids. *Cogn. Psychol.* 23, 483–514
23. Proffitt, D.R. *et al.* (1990) Understanding wheel dynamics. *Cogn. Psychol.* 22, 342–373
24. White, P.A. (2009) Perception of forces exerted by objects in collision events. *Psychol. Rev.* 116, 580–601
25. Tenenbaum, J.B. *et al.* (2011) How to grow a mind: statistics, structure and abstraction. *Science* 331, 1279–1285
26. Battaglia, P.W. *et al.* (2013) Simulation as an engine of physical scene understanding. *Proc. Natl. Acad. Sci.* 110, 18327–18332
27. Hegarty, M. (2004) Mechanical reasoning by mental simulation. *Trends Cogn. Sci.* 8, 280–285
28. Schwartz, D.L. and Black, T. (1999) Inference through imagined actions: knowing by simulated doing. *J. Exp. Psychol. Learn. Mem. Cogn.* 25, 116–136
29. Markman, A.B. and Dietrich, E. (2000) In defense of representation. *Cogn. Psychol.* 40, 138–171
30. Bates, C. *et al.* (2015) Humans predict liquid dynamics using probabilistic simulation. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pp. 172–177, Cognitive Science Society
31. Hamrick, J.B. *et al.* (2016) Inferring mass in complex scenes by mental simulation. *Cognition* 157, 61–76
32. Gerstenberg, T. *et al.* (2015) How, whether, why: causal judgments as counterfactual contrasts. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pp. 782–787, Cognitive Science Society
33. Liang, W. *et al.* (2015) Evaluating human cognition of containing relations with physical simulation. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pp. 1356–1361, Cognitive Science Society
34. Téglás, E. *et al.* (2011) Pure reasoning in 12-month-old infants as probabilistic inference. *Science* 332, 1054–1059
35. Smith, K.A. and Vul, E. (2013) Sources of uncertainty in intuitive physics. *Top. Cogn. Sci.* 5, 185–199
36. Hamrick, J.B. *et al.* (2015) Think again? The amount of mental simulation tracks uncertainty in the outcome. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pp. 866–871, Cognitive Science Society
37. Hespos, S.J. *et al.* (2016) Five-month-old infants have general knowledge of how nonsolid substances behave and interact. *Psychol. Sci.* 27, 2016
38. Kubricht, J.R. *et al.* (2016) Probabilistic simulation predicts human performance on viscous fluid-pouring problem. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pp. 1805–1810, Cognitive Science Society
39. Kawabe, T. *et al.* (2015) Seeing liquids from visual motion. *Vision Res.* 109, 125–138
40. Weiss, Y. and Adelson, E.H. (1998) *Slow and Smooth: A Bayesian Theory for the Combination of Local Motion Signals in Human Vision,* Massachusetts Institute of Technology
41. Lu, H. and Yuille, A.L. *et al.* (2006) Ideal observers for detecting motion: correspondence noise. In *Advances in Neural Information Processing Systems* (18) (Weiss, Y., ed.), In pp. 827–834, MIT Press
42. LeCun, Y. and Bengio, Y. (1995) Convolutional networks for images, speech, and time series. In *The Handbook of Brain Theory and Neural Networks* (Vol. 3361), pp. 255–258, MIT Press
43. Krizhevsky, A. *et al.* (2012) ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (25) (Pereira, F., ed.), In pp. 1097–1105, MIT Press
44. LeCun, Y. *et al.* (2015) Deep learning. *Nature* 521, 436–444
45. Wu, J. *et al.* (2015) Galileo: perceiving physical object properties by integrating a physics engine with deep learning. In *Advances in Neural Information Processing Systems* (28) (Cortes, C., ed.), In pp. 127–135, MIT Press
46. Turing, A.M. (1950) Computing machinery and intelligence. *Mind* 59, 433–460
47. Spelke, E.S. and Kinzler, K.D. (2007) Core knowledge. *Dev. Sci.* 10, 89–96
48. Baillargeon, R. (2002) The acquisition of physical knowledge in infancy: a summary in eight lessons. In *Blackwell Handbook of Childhood Cognitive Development* (Vol. 1) (Goswami, U., ed.), In pp. 46–83, Wiley Blackwell
49. Baillargeon, R. (1994) How do infants learn about the physical world? *Curr. Dir. Psychol. Sci.* 3, 133–140
50. Spelke, E. (1994) Initial knowledge: six suggestions. *Cognition* 50, 431–445
51. Cohen, A.L. (2006) Contributions of invariants, heuristics, and exemplars to the visual perception of relative mass. *J. Exp. Psychol. Hum. Percept. Perform.* 32, 574–598
52. Grzeszczuk, R. *et al.* (1998) Neuroanimator: fast neural network emulation and control of physics-based models. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 9–20, ACM
53. Battaglia, P. *et al.* (2016) Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems* (29) (Lee, D.D., ed.), In pp. 4502–4510, MIT Press
54. Lerer, A. *et al.* (2016) *Learning physical intuition of block towers by example.* arXiv 1603.01312
55. Davis, E. and Marcus, G. (2016) The scope and limits of simulation in automated reasoning. *Artif. Intell.* 233, 60–72
56. Lake, B.M. *et al.* (2016) *Building machines that learn and think like people.* arXiv 1604.00289
57. Caramazza, A. *et al.* (1981) Naive beliefs in 'sophisticated' subjects: misconceptions about trajectories of objects. *Cognition* 9, 117–123
58. McCloskey, M. and Kohl, D. (1983) Naive physics: the curvilinear impetus principle and its role in interactions with moving objects. *J. Exp. Psychol. Learn. Mem. Cogn.* 9, 146–156
59. Runeson, S. (1983) In *On Visual Perception of Dynamic Events* (Vol. 9), Almqvist & Wicksell International