COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science (2018) 1–17

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12596

Sensitivity to the Sampling Process Emerges From the Principle of Efficiency

Julian Jara-Ettinger,[a†] Felix Sun,[b†] Laura Schulz,[b] Joshua B. Tenenbaum[b]

[a]*Department of Psychology, Yale University*
[b]*Department of Brain & Cognitive Sciences, MIT*

Received 28 August 2016; received in revised form 16 January 2018; accepted 18 January 2018

Abstract

Humans can seamlessly infer other people's preferences, based on what they do. Broadly, two types of accounts have been proposed to explain different aspects of this ability. The first account focuses on spatial information: Agents' efficient navigation in space reveals what they like. The second account focuses on statistical information: Uncommon choices reveal stronger preferences. Together, these two lines of research suggest that we have two distinct capacities for inferring preferences. Here we propose that this is not the case, and that spatial-based and statistical-based preference inferences can be explained by the assumption that agents are efficient alone. We show that people's sensitivity to spatial and statistical information when they infer preferences is best predicted by a computational model of the principle of efficiency, and that this model outperforms dual-system models, even when the latter are fit to participant judgments. Our results suggest that, as adults, a unified understanding of agency under the principle of efficiency underlies our ability to infer preferences.

Keywords: Social cognition; Sensitivity to sampling; Principle of efficiency; Theory of mind; Bayesian models of cognition

1. Introduction

As humans, we understand that other people have minds, and we can infer what they know and what they want by watching their behavior. Imagine, for instance, that a man walks toward a cookie jar, opens it, peeks in, and then closes it again. Although we cannot see the inside of the man's mind or of the cookie jar, we nevertheless suspect that

Correspondence should be sent to Julian Jara-Ettinger, Department of Psychology, Yale University, New Haven, CT 06520. E-mail: julian.jara-ettinger@yale.edu

†These authors made an equal contribution to the work.

the man likes cookies, that he planned to eat a cookie, that he believed there were cookies in the cookie jar, and that he was wrong: The cookie jar was empty.

Our ability to infer other people's goals and preferences is at the heart of social cognition. A large body of work suggests that we can infer preferences by relying on two sources of information: how agents navigate in space (spatial information), and the alternatives that they have (statistical information). These capacities, however, have been studied using different paradigms, and they have been explained using different accounts. Our goal in this paper is to explore the prospects for a unified account, showing that both statistical- and spatial-based preference inferences can be explained based on a single principle of efficiency. We present a computational model based on efficiency that can explain classic findings in both paradigms, and we show that it can also explain quantitatively how people trade-off statistical and spatial information about preferences in experimental settings where both sources of information are available.

Efficiency-based reasoning is already the standard way to explain one class of preference inferences: those based on how agents navigate in space. If, for instance, we watch an agent navigate toward an apple and away from an orange, we can infer that she probably prefers apples to oranges, even before she has reached the apple. Our expectation that agents are efficient is so central to how we interpret other people's behavior that, when an agent does not appear to act efficiently, we infer that the agent had incorrect beliefs (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017), that the agent completed a subgoal within the path (Baker, Saxe, & Tenenbaum, 2009), or that the actions themselves are the goal (Schachner & Carey, 2013). These results have been described both formally and informally as evidence for a principle of efficiency (Baker et al., 2009; Csibra, Biró, Koós, & Gergely, 2003; Gergely & Csibra, 2003; Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015; Jara-Ettinger, Schulz, & Tenenbaum, 2015; Jara-Ettinger, Tenenbaum, & Schulz, 2015; Jern, Lucas, & Kemp, 2011; Johnson & Rips, 2015; Lucas et al., 2014), and they trace back to early in infancy. Even 3-month-olds assume that agents take the most efficient path toward their goals, subject to constraints imposed by the environment (Gergely & Csibra, 2003; Liu & Spelke, 2017; Skerry, Carey, & Spelke, 2013).

Standard accounts of how people infer preference from the distribution of the alternatives an agent chooses among do not invoke efficiency, but rather reasoning about statistics and suspicious coincidences in the agent's sampling process. Suppose that an agent can pick a fruit from a bag filled with a hundred apples and just a few oranges. Intuitively, if the agent takes an apple, she does not necessarily prefer apples to oranges. But if she takes one of the very few oranges, then she probably prefers oranges to apples. This ability to draw inferences based on the statistical distribution of the choices, known as sensitivity to the sampling process, has its roots in infancy and it plays a role in how we learn what other people like (Kushnir, Xu, & Wellman, 2010; Wellman, Kushnir, Xu, & Brink, 2016), how we learn about the world (Gweon, Tenenbaum, & Schulz, 2010), and even how we learn the meaning of new words (Xu & Tenenbaum, 2007a,b). Formally, these inferences have been proposed to be guided by the assumption that agents sample objects at random, unless they have a specific preference, in which case they selectively sample from the subset of desirable objects.

Together, these two lines of research suggest that there are two fundamental capacities in social cognition: an ability to infer preferences from spatial information that relies on the principle of efficiency, and an ability to infer preferences from statistical information that relies on sensitivity to the sampling process. But real-world situations do not break down so cleanly. Agents usually combine both spatial and statistical information of potential rewards in their environment, and so should our judgments about their preferences from observing their actions.

One possibility is that when we infer other people's preferences, we combine these two capacities to integrate information about how the agent navigates in space and information about the statistical distribution of choices. In this paper we propose an alternative. Spatial-based and statistical-based preference inferences, and inferences in contexts where both sources of information are at play, may not emerge from two distinct systems of knowledge, but from the principle of efficiency alone. We already noted how preference inferences from navigation in space arise from the principle of efficiency. However, preference inferences from statistical information may also emerge from the same assumption.

To illustrate how sensitivity to the statistical distribution of choices can emerge from the assumption that agents are efficient, consider an agent who can take one piece of fruit from a box filled with apples and oranges (Fig. 1). If the agent has no preference and she acts efficiently, then she should take whichever food is easiest to get. If the fruits in the

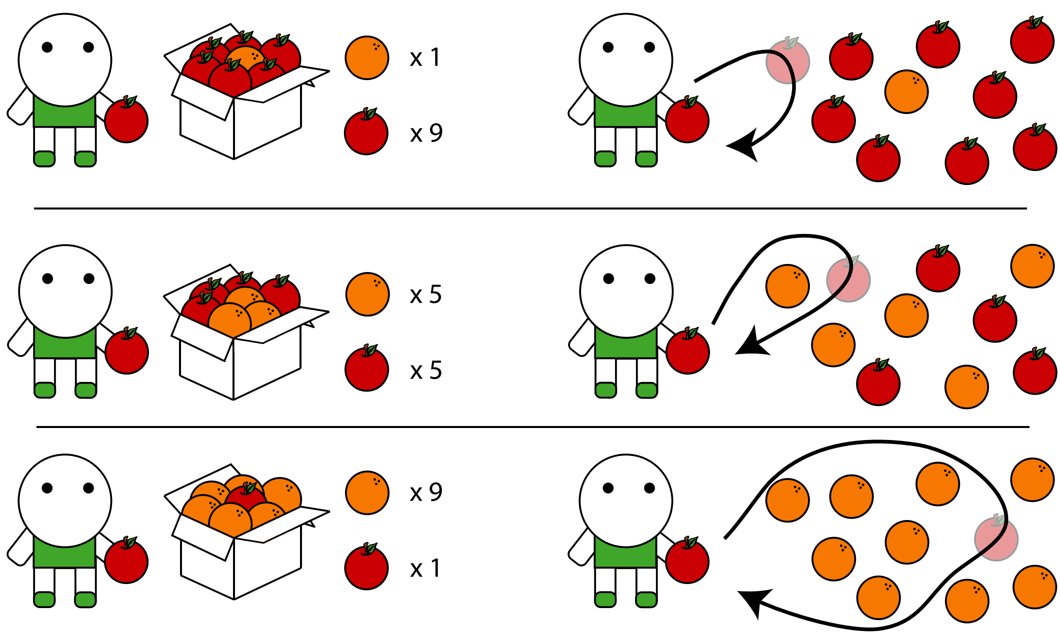

Fig. 1. The left side shows agents' choices in a statistical context, and the right side shows equivalent choices in a spatial context. The rarer a choice is, the more likely that the agent will have to incur additional costs to obtain it.

box are intermixed, then the probability that an apple is easiest to get is given by the proportion of apples in the box (and similarly with oranges). That is, efficient action under no preference produces random sampling (where the randomness is a reflection of which object happens to be easiest to get). In contrast, if the agent prefers apples to oranges, she will have to locate and retrieve one from the box, and this will come at a cost in terms of time, effort, attention, and distance. That is, efficient action under a preference produces selective sampling at a cost. If apples are common, this cost will be negligible. But the less common apples are, the higher the cost associated with retrieving one will be. This is because the rarer apples are, the less likely that there will be an apple conveniently located and easily retrievable (see Fig. 1). Thus, retrieving rare objects suggests that the agent incurred a higher cost. And, if agents are efficient, this additional cost can only be justified if apples have a higher reward relative to oranges. This shows how sensitivity to statistical information may emerge from the principle of efficiency.

Below we formalize this proposal in a way that can be quantitatively tested in behavioral experiments, by implementing a total of five computational models expressing different combinations of efficiency-based, spatial, and statistical reasoning (see Fig. 2). We compare these models against people's judgments in a preference-inference task that combines spatial and statistical information, forcing participants to consider different trade-offs between these factors across different trials. We show that a model based on efficiency alone fits participants' judgments better than models which do not use efficiency, and adding a specific sensitivity to statistical information does not improve fits beyond the efficiency-alone model.

Computational modeling

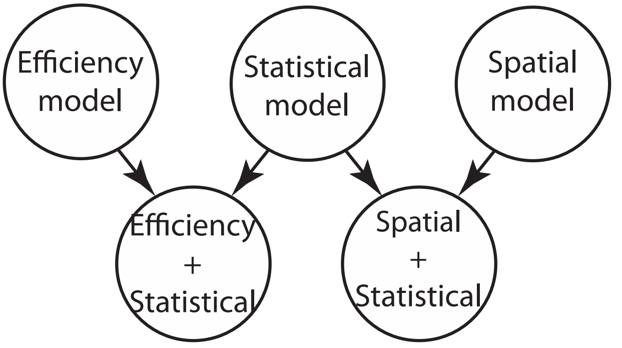

Fig. 2. Schematic of models we compare. Arrows show dependencies between models. The three models at the top are derived from first principles and have no numerical parameters fit to data. The two models on the bottom row are built as linear combinations of the other models. The efficiency model infers preferences through the assumption that agents navigate efficiently. The statistical model infers preferences by assuming that rare choices reveal stronger preferences. The spatial model infers preferences by looking at the spatial location of the agent, without an expectation for efficient navigation in space. The efficiency + statistical and the spatial + statistical model are linear combinations of the other models, with the weight fit to participant judgments.

To be clear, our goal in this work was not to argue that people can attend to both spatial and statistical information in preference inference; that has been well established in previous studies (Gergely & Csibra, 2003; Gweon et al., 2010; Kushnir et al., 2010; Wellman et al., 2016). Neither are we the first to formalize efficiency as the basis for preference inference (Baker et al., 2017). Rather, our goal is to rigorously test if several qualitatively different explanatory principles are necessary to explain preference inferences, or if a single model based on the expectation of efficiency alone can explain how people integrate multiple information sources (specifically, the spatial distribution of rewards and the statistics of different kinds of rewards).

2. Computational modeling

In this section we motivate and describe the five computational models tested in our behavioral study. We begin with a conceptual description of the models and how they relate to each other, and then turn to the formal specifics of each model.

The first model infers preferences through the expectation that agents are efficient (efficiency model). This model captures the principle of efficiency from Gergely and Csibra (2003) with an implementation based on Baker et al. (2009). The second model infers preferences through sensitivity to the statistical information (statistical model). This model captures sensitivity to the sampling process based on Gweon et al. (2010). These two models were derived from first principles and had no parameters fit to human data. If our account is correct, then (a) the efficiency model should fit participant judgments with high accuracy (giving evidence that the expectation for efficiency alone explains people's sensitivity to spatial and statistical information), and (b) the statistical model should broadly fit participant judgments (giving evidence that participants pick up statistical information in our task), but it should fail to explain fine-grained variance (giving evidence that people are not inferring preference based on the statistical information alone, independent of the spatial information).

Our third computational model is designed to test if people infer preferences by combining their understanding that agents act efficiently with their sensitivity to the sampling process (called the efficiency + statistical model). This model generated predictions through a linear combination of the efficiency and the statistical model's predictions, with the weights fit to participant judgments. If our account is correct and people's preference inferences are driven by the principle of efficiency alone, then this compound model should not outperform the efficiency model, even though it has the advantage of being fit to participant data. If, instead, people's inferences are generated by a combination of the two systems of knowledge—the spatial-based and the statistical-based ones—then the compound model should outperform both the efficiency and the statistical model, showing that a combination of both systems are necessary for explaining participant judgments.

As reviewed above, there is substantial evidence that action-understanding is at least in part guided by the expectation that agents navigate efficiently (Baker et al., 2017; Gergely & Csibra, 2003; see Jara-Ettinger et al., 2016, for review). It is possible, however, that inferences that appear to be guided by the principle of efficiency are instead approximated by some combination of a statistical and a spatial model, neither of which includes an explicit expectation for efficient action. If so, sensitivity to the sampling process may one of the components of this larger system. To explore this idea, we designed a fourth model that infers preferences based on the agent's locations in space, without the expectation for efficient behavior (spatial model), and we combined these two models (statistical and spatial model) to build a fifth model (called the spatial + statistical model) that combines the predictions of these two models through a linear combination fit to participant judgments. If our account is correct and preference inferences rely on the principle of efficiency, rather than on a combination of simpler models that approximate the efficiency model, then we expect the following. First, the spatial model should broadly capture participant judgments (giving evidence that our proposal is a plausible alternative), but the spatial + statistical model should perform worse than the efficiency model (giving evidence that preference-inferences are indeed driven by the principle of efficiency, and not by some approximation through two separate systems of knowledge). Thus, together, these five computational models enable us to test if the expectation for efficiency is necessary and sufficient for explaining people's sensitivity to spatial and statistical information.

2.1. Efficiency model

The efficiency model assumes that agents act to minimize costs when seeking rewards. As in past models of the principle of efficiency (Baker et al., 2009, 2017), we rely on frameworks developed in the AI and robotics literature to derive the sequence of actions that an agent should take. Namely, we use a Markov decision process (MDP; Sutton & Barto, 1998).

In MDPs, the world is represented as a set of states S, each with a reward determined by the reward function $R : S \rightarrow R$. Intuitively, the reward function corresponds to the agent's desires. In each time step, the agent takes an action $a \in A$ and changes the state from s_0 to s' with probability $T_a(s_0, s')$ (where T stands for transition). In our model, we set the transitions to move the agent deterministically (i.e., taking the action "up" always moves the agent one cell up, unless it reaches a map limit). This corresponds to an assumption that agents' actions are usually intentional. Using the MDP framework, it is possible to compute the exact sequence of actions that maximizes any reward function. That is, the set of actions that an agent ought to take if they act efficiently. Here, however, we use an expectation for approximate efficiency, rather than exact efficiency. This corresponds to the expectation that agents tend to act efficiently, but that agents also commit minor errors as they navigate. Consistent with past work (e.g., Baker et al., 2017), we do this by applying the soft-max function to the calculation of the optimal policy, the probability that an agent takes action a in state s is given by

$$P(a|s) = \frac{\exp(V^*(s,a)/\tau)}{\sum_{a' \in A} \exp\left(\frac{V^*(s,a')}{\tau}\right)} \tag{1}$$

where $V^*(s,a)$, the optimal value of taking action a in state s, is given by Bellman's equation

$$V^*(s,a) = R(s,a) + \max_{a' \in A} \gamma \sum_{s' \in S} T_a(s,s') V^*(s',a'). \tag{2}$$

The parameter $\tau \in (0, \infty)$ determines the strength of the expectation for efficiency. As $\tau \to 0$, the agent is expected to act optimally, and as $\tau \to \infty$, the agent is expected to move randomly. $\gamma \in (0,1)$ determines the future discount and is necessary for the planner's convergence (Sutton & Barto, 1998).

Using Markov decision processes (MDPs) to model efficient behavior, goal inference can be formalized as the problem of finding the reward function that can produce the observed actions. Intuitively, this corresponds to making sense of other people's behavior through an intuitive theory of how they would act under different circumstances. Versions of these problems, known as inverse reinforcement learning (Ng & Russell, 2000), inverse optimal control (Dvijotham & Todorov, 2010), or inverse planning (Baker et al. 2017), can be solved through Bayesian inference over the generative model. Under this framework, the posterior probability that an agent has the unobservable reward function R given a sequence of n actions $\vec{a} = (a_1, \ldots, a_n)$ is given by Bayes' rule:

$$p(R|\vec{a}) \propto l(\vec{a}|R)p(R) \tag{3}$$

where the likelihood of the actions $\vec{a}$ given the reward function R is determined by Eq. 1.

In simple two-dimensional grid worlds (which capture enough information for humans to make rich social inferences; Heider & Simmel, 1944), goal inference through inverse planning matches human performance with quantitative precision (Baker et al., 2009, 2017; Ullman et al., 2009). Nonetheless, these models have only been used in paradigms that manipulate spatial information but do not integrate information about sampling behaviors. In our implementation, we fixed the future-discount to $\gamma = 0.999$, the rationality to $\tau = 0.01$, and a constant cost per action set to 0.01. These values were not fit to our data and were instead set to minimize their impact as done on past work (Baker et al., 2009, 2017; Jara-Ettinger, Schulz, et al., 2015).[1]

2.2. Statistical model

The statistical model is inspired by proposals for how people infer preferences by relying on statistical information (Gweon et al., 2010; Xu & Tenenbaum, 2007a,b). These models were formulated in simpler domains than the one we test in our experiment so we extended them to fit our experimental design. Classically, these models are presented as direct mappings from actions to preferences. Here we formalize preferences as relative rewards in order to make the relation between models clearer (see also Lucas et al., 2014, for a utility-based formulation of these kinds of sampling models). Our sampling model assumes that agents decide whether to collect an object of that category based on its reward (usually called the preference). Specifically, if there are m objects of n categories, the agent will consider an object from category k with probability $N(k)/m$ where $N(k)$ is the number of objects in category k, and she will decide to collect it with probability

$$r_k \Big/ \sum_{i=1}^{n} r_i \tag{4}$$

Together, these assumptions imply that the strength of the agent's preference (formalized as the relative rewards) and the availability of the preferred objects influence the agent's choice. When the agent has no preference, then the statistical availability alone guides the agent's choices, producing random sampling. When the agent has an overwhelming preference, then this preference trumps the statistical availability, producing selective sampling.

Given the generative model of how agents choose what to collect, we use Bayesian inference to recover the agent's preferences given her choices. Specifically, because in our experiment we use two types of objects (see Section 3.1), we use Bayes' rule to estimate the relative magnitude of one reward type over the other (with 0 indicating that the first category contains all the rewards, 0.5 indicating that both categories are equally rewarding, and 1 indicating that the second category contains all the rewards), using a uniform prior.

2.3. Spatial model

Finally, we implemented a simplified preference-inference model that infers agents' preferences by considering the positions in space the agent moves toward, while ignoring the statistical properties of the objects or the agent's relative efficiency. This model assumes that the probability that an agent will be in a location with an object is given by the object's relative reward (using the same relative reward equation from the statistical model; see Eq. 4). As such, this model only infers preferences based on where the agent goes to, disregarding both the distance the agent had to travel or the statistical distribution of the choices. As in the past two models, preferences are inferred using Bayes' rule.

2.4. Model combinations

Using the three main models, we implemented two model combinations (Fig. 2). The first compound model is designed to enable us to test if the efficiency model alone captures spatial and statistical sensitivity. This compound model linearly combines the predictions from the efficiency and the statistical model (called the efficiency + statistical model). If the efficiency model alone already captures people's sensitivity to the statistical information, then the efficiency + statistical model should not outperform it. The second compound model is designed to explore the possibility that preference-inferences do not rely on any expectation for efficient action and are instead driven by an integration of our spatial and our statistical models. This second compound model linearly combines the predictions from the spatial and the statistical model (called the spatial + statistical model). If the efficiency model underlies people's sensitivity to spatial and statistical information, it should outperform this compound model. In contrast to the main three models, where all numerical parameters were fixed prior to data collection, the parameters in these two compound models were estimated through a linear regression fit to participant judgments (see Supporting Information for details).

3. Experiment

To test our models, we designed a simple task where participants watched a miner collect minerals in mines with variable distributions of minerals.

3.1. Stimuli

Fig. 3 shows examples of the stimuli. Each stimulus consisted of an animated display of an agent (the miner) entering a mine (a 12×12 grid world) and collecting green and/or red minerals. Each map contained 24 minerals in the same locations (which were chosen at random and kept constant across stimuli), but the proportion and the distribution of these minerals varied. The proportion of minerals varied according to three levels: more green than red (20 green and 4 red), more red than green (4 green and 20 red), or an equal number of each (12 of each). The distributions of these minerals varied according to three levels: red minerals closer, green minerals closer, or all minerals intermixed. This generated a total of nine different maps (see Fig. 3 for examples; see Supporting Information for all maps). By varying the proportion of the objects, we can test how statistical information influences preference inferences; by varying the location of the objects, we can test how spatial information influences preference inferences.

The miner's paths were obtained by computing the shortest path an agent would need to take to collect all minerals of the kind that is farther away from the entrance (called the "strong preference" path types; see, e.g., panels A, B, D, G, and I in Fig. 3) or to collect the closest minerals (called the "no preference" path types; which consisted either of a selection of both red and green minerals, or all minerals of the same color, depending on which minerals were closer to the map entrance. See panels C, E, F, and H in Fig. 3 for examples). These paths were generated in accordance to three conditions. In the first condition, the miner collected one mineral and exited the mine (see, e.g., panels A–C in Fig. 3). In the second condition, the miner collected three minerals in a single trip and then exited the mine (see, e.g., panels D–F in Fig. 3). And in the last condition the miner collected three minerals, but had to return to the mine's exit after collecting each object (see, e.g., panels G–I in Fig. 3). Thus, the first and second conditions test how the amount of data an observer receives influences observers' inferences, and the second and third conditions together test how the costs of collecting the minerals

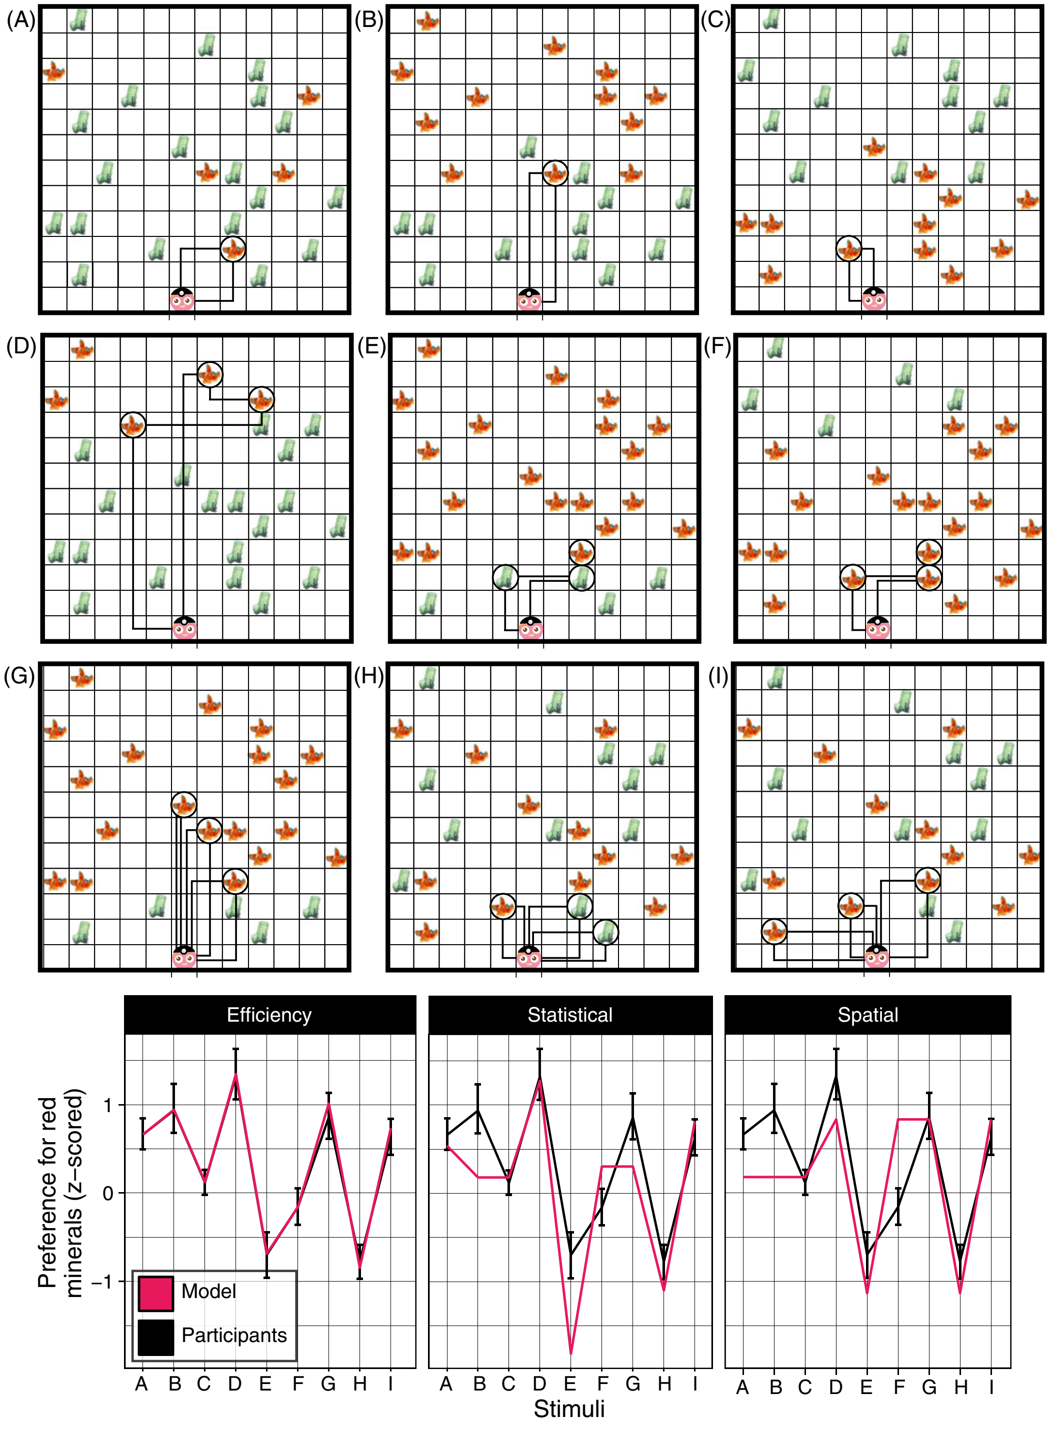
(A)
(B)
(C)
(D)
(E)
(F)
(G)
(H)
(I)
Efficiency
Statistical
Spatial
Preference for red minerals (z–scored)
1
0
–1
Model
Participants
A B C D E F G H I
Stimuli

Fig. 3. Nine example trials along with average human responses (z-scored) and model predictions. The top row shows three examples of the condition where the agent only collects one mineral; the middle shows examples of the condition where the agent collects three minerals in one trip; and the bottom row shows examples of the condition where the agent collects three minerals in three trips.

influence observers' inferences. The combination of the two path types ("strong preference" and "no preference") with the nine maps produced a total of 18 test paths per condition.

3.2. Participants

Ninety U.S. residents (as determined by their IP address) were recruited and tested through Amazon's Mechanical Turk platform (M_{age} = 33 years; range = 20–59 years). No gender or demographic information was collected on participants. No participants were excluded from the study.

3.3. Procedure

Participants were randomly assigned to the one mineral condition, to the three minerals in one trip condition, or to the three minerals in three trips condition (N = 30 participants per condition; each condition consisting or 18 trials). To ensure participant's judgments were as independent as possible across trials, we explained to participants that the miners in the task were mining for the purpose of selling the minerals, and that the relative value of the minerals changed each day. Thus, the miners' relative preference directly reflected the relative market value of the minerals. This cover study enabled us to (a) avoid potential task demands of people jointly inferring objective (market value) and subjective (preference) properties of the minerals, and (b) to avoid participants from attempting to learn spurious correlations across trials (by potentially assuming that people often have similar preferences). Participants next completed a questionnaire with three questions to ensure that they understood the task (see Supporting information for details). Only participants who responded all questions correctly were given access to the experiment. Participants who made at least one error were redirected to the beginning of the tutorial and were allowed to reread the tutorial and try again.

In the test stage, participants saw an animated display of the miner collecting the minerals in each trial, and they were asked to respond four questions. The first two questions were multiple-choice control questions asking about the proportion and distribution of the minerals in the map (see Supporting Information for an example of the test screen). Participants who answered these questions incorrectly were not allowed to proceed to the next trial and were asked to re-examine the stimulus. The third and fourth questions were the critical test questions. The third question asked participants to rate the miner's preference using a continuous slider that ranged from "Red is much more valuable" (coded as a 0) to "Green is much more valuable" (coded as a 1). The final question asked participants to rate how confident they were in their inference using a continuous slider that ranged from “Not at all” (maximum uncertainty; coded as a 1) to “Extremely confident” (minimal uncertainty; coded as a 0).

4. Results

Fig. 4 shows the model predictions plotted against average participant responses to the first test question. From the main three models (efficiency, statistical, and spatial), the efficiency model had the highest correlation (r = .97) between its predictions and participant responses, and this correlation was significantly higher than the correlation of the statistical (correlation difference = 0.16; 95% CI = (0.05, 0.22); see bottom-right panel in Fig. 4) and the spatial model (correlation difference = 0.15; 95% CI = (0.05, 0.21); see Fig. 4).

Fig. 3 reveals why the efficiency model fit participant judgments better than the statistical and the spatial model. The statistical model makes a weaker inference in display B compared to display A because the proportion of red minerals is higher. But it does not account for the fact that the agent travelled farther into the mine to collect it, as participants and the efficiency model do. Similarly, the statistical model infers a strong dispreference for red minerals in display E and does not account for the fact that the green minerals happened to be closer. The spatial model’s failure is evident when considering displays D, F, G, and I. In these displays, the agent always selects three locations with red minerals, and as such, the spatial model makes the same inference. Although these inferences match participant judgments qualitatively, only the efficiency model shows the sensitivity to the cost the agent incurred (both in terms of distance travelled and number of trips the agent had to follow) and sensitivity to the statistical distribution of the objects.

These results so far confirm the first two predictions outlined in the introduction: the statistical and the spatial model predict participant judgments at a coarse level (providing evidence that participants were indeed sensitive to the statistical and the spatial information in our task), and only the efficiency model predicts participant judgments at a fine-grained level (providing evidence that participants’ sensitivity to spatial and statistical information can be explained through the principle of efficiency alone).

Next, we compared the two compound models against the efficiency model. As Fig. 4 shows, the efficiency model performed equally well as the model that combined the efficiency and statistical models, even though the later model was fit to parameter data (95% CI over correlation difference: −0.02–0.02), whereas the former was not. This confirms our prediction that the efficiency model alone explains people’s sensitivity to the sampling process, and that integrating a second system of statistical sensitivity does not improve model fit. The combination of the spatial and statistical models (Spatial + Statistical) performed reliably worse than the efficiency model (95% CI over correlation difference: 0.04–0.19). This suggests that the principle of efficiency underlies participants’ sensitivity to spatial and statistical information, and that it is not instead approximated by

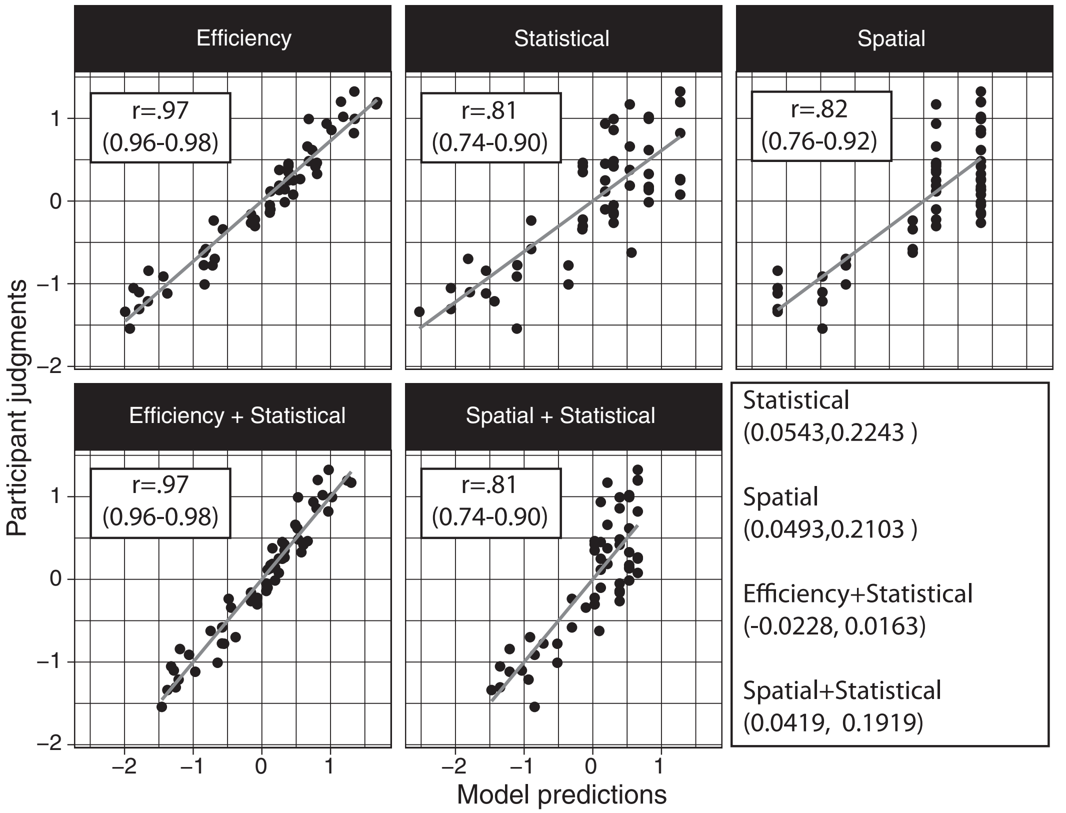

Fig. 4. Experiment results. In each plot, each black dot represents a stimulus. The *x*-axis shows the model's prediction (*z*-scored) and the *y*-axis shows average participant judgments (*z*-scored within each participant and averaged). Each plot presents the model's correlation along with the 95% confidence interval. The bottom right panel presents the 95% confidence intervals of the difference between the efficiency model and each of the alternative models.

two separate systems, neither of which explicitly includes an expectation for efficient behavior.

Finally, we compared the standard deviation of each model's posterior distribution against participants' confidence judgments. We do not expect the standard deviation to directly determine participants' confidence, because we have no principled reason to believe that standard deviation (as opposed to other potential ways to quantify uncertainty) should capture participant uncertainty at a cognitive level. Nonetheless, it is a useful metric to characterize how each model's uncertainty changes as a function of parametric variations of spatial and statistical information. If participants' inferences are similar to those of one of our models, then their confidence judgments should correlate with the model's standard deviation. The results were consistent with our main analyses. The efficiency model had the best correlation against participants' confidence judgments (r = .65; 95% CI: 0.46–0.78). The statistical model and the spatial models had a reliably lower correlations (r = .28; 95% CI: 0.01–0.51 for the statistical model, and r = .04; 95% CI: −0.23–0.30 for the spatial model). Finally, the combination of the efficiency and statistical model did not fit confidence judgments better than the efficiency model alone (r = .65; 95% CI: 0.46–0.78), and the spatial + statistical model was also reliably lower (r = .46; 95 CI: 0.10–0.57).

5. Discussion

Here we reviewed past research that shows that, from early in life, humans can infer preferences based on spatial information (how agents navigate through space) and on statistical information (the rarity of their choices). Past work has usually treated these two types of inferences as separate mechanisms. Here we proposed that, in adults, sensitivity to spatial and to statistical information arises from the principle of efficiency alone. Consistent with past literature, our empirical results confirm that adults were sensitive to both the spatial and the statistical information when inferring preferences. Our modeling analyses show that the efficiency model (a) captures human participants with high precision, (b) that it outperforms alternative accounts proposed in the literature, and (c) that more complex models built on top of it do not outperform the efficiency model, even when fit to participant data. Together, these findings suggest that people infer preferences through a unifying theory of agency with the expectation for efficient behavior at its heart.

Critically, all the accounts we presented were qualitatively consistent with our data, and an analysis based on the qualitative patterns would have been unsuccessful in distinguishing their performance. Thus, implementing formal computational models was critical for generating precise predictions and assessing whether they explained variation in human judgments in a fine-grained manner. Altogether, our results show that the principle efficiency explains why and how humans rely on spatial and statistical information when inferring preferences.

For each map in our stimuli (see Supporting Information and Stimuli section) we designed two kinds of paths, one where the agent navigated to the objects that were costlier to collect—called "strong preference" paths—and paths where the agent navigated to the objects that were closer—called "no preference" paths. In the "strong preference" paths, the agent must have had a strong preference (motivating the agent to go past the more convenient objects), and both our model and participants inferred this. In the "no preference" paths, the closest objects were either all of the same color or a combination of red and green minerals. When the closest minerals were a combination of red and green minerals, the agent collected a combination of these (see, e.g., panels E and H in Fig. 3). In this case, the agent probably had a weak or no preference (otherwise she would have selectively collected one type of mineral). Both participants and our model captured this. In contrast, in the "no preference" paths where the closest objects were all of the same kind (see, e.g., panels C and F in Fig. 3), having a strong preference for the nearby objects produces the same behavior than having no preference at all: In both cases, the agent will collect the objects that are closest. As such, these paths are judged by the model and our participants as not revealing a strong preference, but they carry less confidence. This kind of asymmetry, where costlier choices reveal preference more clearly has been explored in the past (see, e.g., Experiment 1 in Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015) and is a signature in both statistical contexts (rare choices reveal preferences more clearly) and in spatial contexts (far away choices reveal preference more clearly). The fact that this asymmetry appears in both spatial and statistical contexts, and that an efficiency model alone explains them by appealing to costs (costlier choices reveal preferences more clearly) provides further qualitative evidence for our account.

Our work is consistent with other lines of work exploring how we infer preferences. In related work, Lucas et al. (2014) showed that some formulation of inverse-utility reasoning (which, at a high level, is conceptually similar to our efficiency model) explains preference inferences in a wide range of tasks that even children and toddlers succeed at. Although Lucas et al. (2014) established how utility-based models can explain preference inferences, it did not explore the role of efficient navigation in space in action-understanding. Our work complements the work from Lucas et al. (2014) by expanding this formulation to handle preference inferences in spatial domains and to show how the spatial navigation formulation can naturally give rise to the dynamics of sensitivity to statistical information.

Empirical results show that the ability to infer preferences from both spatial information and statistical information arise in early childhood (Gergely & Csibra, 2003; Gweon et al., 2010). However, these sources of information have been studied separately, and different accounts have been proposed to explain how we draw these inferences. Here we emphasized the developmental literature not because our adult studies have direct implications for how to interpret those studies, but rather because they are critical to establishing that sensitivity to spatial and statistical information is a constitutive part of social cognition that underlies the most basic ways we reason about goal-directed action, and it is already at work by our first year of life. Our finding that inferences from efficient navigation in space and from the rarity of the choices are unified in adults opens questions about the developmental origins. A wealth of evidence already suggests that even infants have a principle of efficiency (Gergely & Csibra, 2003; Liu & Spelke, 2017), and so on. Although future work is necessary, this opens the possibility that evidence of sensitivity to statistical information may be tapping into the same fundamental understanding and providing evidence that the principle of efficiency is even more powerful than previously believed.

Acknowledgments

This material is based upon work supported by the Center for Brains, Minds, and Machines (CBMM), funded by NSF-STC award CCF-1231216, and by the Simons Center for the Social Brain.

Note

1. Note that our model evaluations are done using z-scored predictions, minimizing the influence of these parameters even further.

References

Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, *1*(4), 0064.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113* (3), 329–349.

Csibra, G., Bıró, S., Koós, O., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science*, *27*(1), 111–133.

Dvijotham, K., & Todorov, E. (2010). Inverse optimal control with linearly-solvable MDPs. In J. Fünkranz and T. Joachims (Eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML-10)* (pp. 335–342). Madison, WI: Omnipress.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naıve theory of rational action. *Trends in Cognitive Sciences*, *7*, 287–292.

Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*(20), 9066–9071.

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, *57*(2), 243–259.

Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition*, *140*, 14–23.

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: Computational principles underlying commonsense psychology. *Trends in cognitive sciences*, *20*(8), 589–604.

Jara-Ettinger, J., Schulz, L. E., & Tenenbaum, J. B. (2015). The naive utility calculus: Joint inferences about the costs and rewards of actions. Proceedings of the 37th Annual Conference of the Cognitive Science Society, Austin, TX.

Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not so innocent toddlers' inferences about costs and culpability. *Psychological Science*, *26*(5), 633–640.

Jern, A., Lucas, C. G., & Kemp, C. (2011). Evaluating the inverse decision-making approach to preference learning. In J. Shawe-Taylor, R. S. Zemel, P. Bartlett, F. Pereira, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 24, pp. 2276–2284). La Jolla, CA: The NIPS Foundation.

Johnson, S. G., & Rips, L. J. (2015). Do the right thing: The assumption of optimality in lay decision theory and causal judgment. *Cognitive Psychology*, *77*, 42–76.

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*(8), 1134–1140.

Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition*, *160*, 35–42.

Lucas, C. G., Griffiths, T. L., Fei, X., Fawcett, C., Gopnik, A., Kushnir, T., Markson, L., & Jane, H. (2014). The child as econometrician: A rational model of preference understanding in children. *PLoS ONE*, *9*(3), e92160.

Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In P. Langley (Ed.), *ICML* (pp. 663–670). San Francisco: Morgan Kaufmann.

Schachner, A., & Carey, S. (2013). Reasoning about "irrational" actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*, *129*, 309–327.

Skerry, A. E., Carey, S. E., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in prereaching infants. *Proceedings of the National Academy of Sciences*, *110*(46).

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum, J. (2009). Help *or* hinder*:* Bayesian models *of* social goal inference. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems* (Vol. 22, pp. 1874–1882). La Jolla, CA: The NIPS Foundation.

Wellman, H. M., Kushnir, T., Xu, F., & Brink, K. (2016). Infants use statistical sampling to understand the psychological world. *Infancy*, *21*, 668–676.

Xu, F., & Tenenbaum, J. B. (2007a). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245.

Xu, F., & Tenenbaum, J. B. (2007b). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, *10*(3), 288–297.

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:

Appendix S1. Experiment details and supplemental analyses.