Clustering and the efficient use of cognitive resources

Ishita Dasgupta

DeepMind New York City

Thomas L. Griffiths

Princeton University

Author Note

Abstract

A central component of human intelligence is the ability to make abstractions, to gloss over some details in favor of drawing out higher-order structure. Clustering stimuli together is a classic example of this. However, the crucial question remains of how one *should* make these abstractions – what details to retain and what to throw away? How many clusters to form? We provide an analysis of how a rational agent with limited cognitive resources should approach this problem, considering not only how well a clustering fits the data but also by how 'complex' it is, i.e. how cognitively expensive it is to represent. We show that the solution to this problem provides a way to reinterpret a wide range of psychological models that are based on principles from non-parametric Bayesian statistics. In particular, we show that the Chinese Restaurant Process prior, ubiquitous in models of human and animal clustering behavior, can be interpreted as minimizing an intuitive formulation of representational complexity.

Keywords: Bayesian inference; resource rationality; information theory; probabilistic numerics

Clustering and the efficient use of cognitive resources

The only enduring aspect of our environment is that nothing stays the same. We never have *exactly* the same experience twice. As a consequence, the human mind has to form abstractions, clustering these experiences together in a way that supports generalization. Psychological models have applied this lens to phenomena as diverse as categorization (Anderson, 1991), feature learning (Griffiths & Austerweil, 2008), theory formation (Kemp, Tenenbaum, Niyogi, & Griffiths, 2010), classical conditioning (Gershman, Blei, & Niv, 2010), and word segmentation (Goldwater, Griffiths, & Johnson, 2009). A key problem that arises in each of these models is knowing when to generate a new cluster – when an object, stimulus, or word is genuinely of a kind that has never been seen before.

Deciding when to form a new cluster involves making a trade-off between the complexity of the underlying representation and how well it describes the environment. Grouping all experiences into a single cluster where they are represented by some abstract summary statistics is maximally simple, but at the cost of losing a significant amount of detail. Having a separate cluster for each experience accurately captures the nuances of those experiences, but is maximally complex. So how *should* we form clusters?

In this paper, we address this question in the spirit of rational analysis (Anderson, 1991), asking how it might be solved by an ideal agent. More precisely, we engage in resource rational analysis (Gershman, Horvitz, & Tenenbaum, 2015; Griffiths, Lieder, & Goodman, 2015; Lieder & Griffiths, 2019), since our analysis focuses on the question of how that agent might make the best use of limited cognitive resources. Formalizing the complexity of a clustering in information-theoretic terms, we derive an optimal distribution over clusterings.

This analysis yields a surprising result: our optimal solution has the same properties as the distribution over clusterings assumed in *all* of the psychological models mentioned above. These models use a distribution over clusterings originally introduced in psychology

by Anderson (1991) in his rational model of categorization. This distribution was independently discovered in non-parametric Bayesian statistics (Aldous, 1985; Hjort, Holmes, Müller, & Walker, 2010), where it is known as the Chinese restaurant process (CRP).[1]

While the CRP has a number of attractive mathematical properties, none of these properties justify why it might make sense as an assumption in psychological models. Our analysis shows that CRP-like distributions can arise from an effort to minimize representational costs, i.e. that this distribution is normative under an assumption of resource-rationality. In particular, we show that minimizing the number of bits required to store an object-category mapping (viz. reducing the complexity of that mapping) can give CRP-like behavior. These results provide the first process-level explanation of why this kind of clustering behavior might be a reasonable assumption in psychological models.

The plan of the paper is as follows. In the next section we provide a more detailed introduction to the Chinese restaurant process. We then turn to our analysis of optimal clustering under resource constraints. We derive our key results mathematically and present simulations that verify our analysis. We conclude with a discussion of the implications of these results for developing models of human cognition.

Background: The Chinese restaurant process

As mentioned above, one of the challenges involved in clustering a set of experiences is deciding how many clusters there should be. Researchers in nonparametric Bayesian statistics developed an innovative strategy for solving this problem: rather than specifying a particular number of clusters, instead assume that there could exist an infinite number of clusters of which only a finite number have been observed so far. The problem of determining the number of clusters then becomes a matter of inferring how many clusters

---

[1] The name of this process comes from its creators imagining a large restaurant that seats multiple parties at communal tables, with people joining tables based on their current popularity – a phenomenon that could apparently be observed in San Francisco's Chinatown. The tables are clusters and the people the experiences being clustered.

may have been observed, which can be solved by applying Bayes' rule.

Pursuing this approach requires identifying a prior distribution over clusterings that remains well-defined regardless of how many experiences need to be clustered. A common way to achieve this goal is to assume that the prior probability an item belongs to a cluster follows a distribution known as the Chinese Restaurant process (CRP; Aldous, 1985). Under this distribution, the probability of belonging to an existing cluster is proportional to the number of objects already in that cluster, while the probability of a new cluster is proportional to the value of a parameter α.

The CRP has various desirable properties that make it a sensible choice. In the infinite space of possible clusterings, it favors assigning data to a small number of clusters. The expected number of clusters grows slowly as the number of experiences being clustered increases. In particular, the CRP displays 'preferential attachment' or a 'rich-get-richer' dynamic, where a cluster with a large number of members is more likely to grow further. The resulting distributions of cluster sizes ('scale-free' distributions, where the size of clusters decays as a power law) have been shown to be prevalent across several other domains (Adamic & Huberman, 2002; Barabási & Albert, 1999; Mandelbrot, 1960; Rosen & Resnick, 1980).

Another practical reason for the success of the CRP is that it is agnostic to the order of data presentation (i.e. exchangeable Aldous, 1985) – changing the order of presentation of experiences does not change the probability of their cluster memberships. This makes Bayesian inference more tractable, as it is easy to compute conditional distributions that are required for standard inference algorithms (see, e.g., Gershman & Blei, 2012).

However, none of these properties of the CRP provide a cognitive justification for why human minds might use this particular prior distribution for clustering. In its first use in psychology, Anderson (1991) derived the CRP from the assumption that any two objects must have the same fixed prior probability of being in the same cluster. This is related to exchangeability. While this might be a useful property to have, it is odd to think of it as a

reason that we use this particular prior – it is not normative. In other words, this kind of exchangeability (as well as the various properties of the CRP) might be emergent properties of a cognitively justifiable process that gives rise to CRP-like behavior. This is the possibility we pursue in the remainder of the paper.

Resource-rational clustering

As noted above, the CRP has been used to model clustering problems that arise in a variety of domains. Following Anderson's (1991) original application we will focus on the case where the agent seeks to organize a set of objects into clusters to support their categorization (see Figure 1). We formalize this problem as follows. We have a total of N objects. Since we are concerned with examining the *prior* over clusterings (i.e., how each object should be assigned to a cluster in the absence of any specific features), we assume that these objects do not have any distinguishing features except for their index $i \in [0, N]$. The goal is to organize these objects into clusters. We don't know a priori how many clusters the objects will be sorted into, but they will certainly be no more by the number of objects N. We therefore need to learn a mapping π from object o_j for $j \in [0, N)$ to cluster c_i for $i \in [0, N)$.

For an agent with finite cognitive resources, it will be important to represent these objects in as simple a way as possible while allowing for the potential differences between them. We will derive a prior based on this idea of learning 'simpler' mappings π, and show that this simplicity prior corresponds closely to the CRP. But first, how do we measure simplicity?

Measuring simplicity or complexity

We will use the representational cost of a mapping π, as a measure of its simplicity or complexity: a representationally more expensive mapping is more 'complex'. We follow previous work (e.g., Bhui & Gershman, 2018; Gottwald & Braun, 2019; Olshausen & Field, 1996; Ortega, Braun, Dyer, Kim, & Tishby, 2015; Todorov, 2009; Zenon, Solopchuk, &
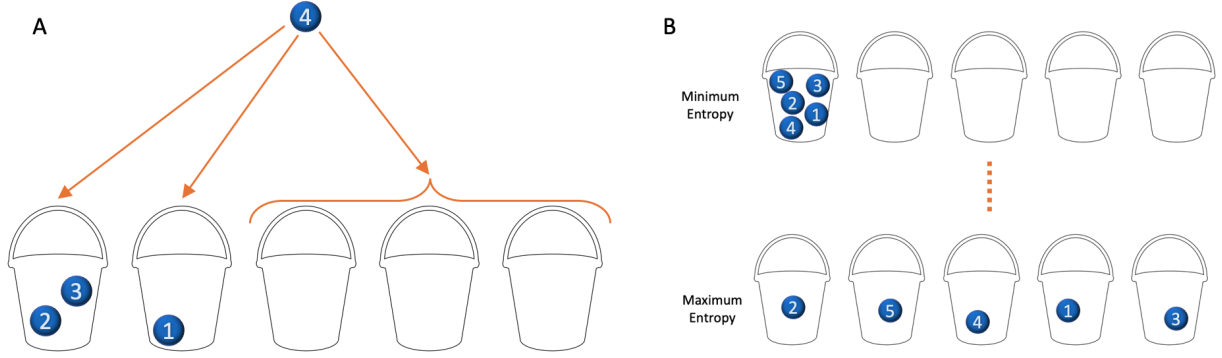
Figure 1. **Schematic illustration of the clustering problem.** In this schematic, we represent each object as a blue ball and each cluster as a bucket. Each object has no features except its unique index. A) We want to cluster object 4, conditioned on how objects 1, 2, and 3 were clustered. There are three possibilities (indicated by the orange arrows): assign to bucket 1 which contains two objects already, assign to bucket 2 with one object, or start a new cluster by assigning to one of the empty buckets. In this work, the prior over clusterings specifies these conditional probabilities. B) By iterating these conditional probabilities, we can derive a probability distribution over the range of possible final clusterings of all objects (in this case 5 objects in total). These final clusterings are represented here, as well as how they differ in the entropy of the marginal distribution over clusters. A prior over clusterings specifies a distribution over these different solutions.

Pezzulo, 2019), in borrowing from information theory to define this representational cost. We make this more precise below.

We first compute an intermediate quantity, the marginal distribution over categories given a mapping π:

$$P_\pi(c_i) = \frac{\sum_j^N \mathbb{I}[\pi(o_j) == c_i]}{N} \tag{1}$$

Each mapping π gives a probability distribution over categories. We then define simplicity or complexity of this probability distribution. What makes one distribution over clusters more or less complex than another?

The entropy of the distribution can act as a measure of its representational cost and

thereby of its complexity. It is given by:

$$\mathcal{H}(\pi) = -\sum_{i}^{K} P_\pi(c_i)\log P_\pi(c_i) \qquad (2)$$

The information-theoretic interpretation of the entropy of a distribution is that it measures the average number of bits $(0 - 1$ coin flips) required to convey an object sampled from that distribution, under the most efficient code possible. The number of bits required for c_i, or the length of its 'codeword', is $\log P_\pi(c_i)$ (under the most efficient code; Shannon, 1948). Weighting this codeword length by the probability of each token gives the entropy of the distribution. Intuitively, this measures how difficult it is (i.e. how many bits of information are required) to convey which object is sampled, when randomly sampling objects from the given distribution, to an observer that knows the distribution but doesn't know which specific object was sampled. A representationally 'expensive' or 'complex' distribution is one that requires more such bits.

In using entropy as a measure of representational complexity, we are following previous work in both psychology and neuroscience. Work on planning and sequential decision-making has used entropy as a measure of representational cost (Todorov, 2009), and other work has suggested that a tendency to minimize this information-theoretic cost is what characterizes bounded-rational behavior in agents with limited resources (Olshausen & Field, 1996; Ortega et al., 2015). This tendency has been empirically validated, and used to model neural representations (Laughlin, 1981) as well as human behavior in high-level cognitive tasks (Bhui & Gershman, 2018).

How does this measure map onto our intuitions in this domain? The lowest entropy distribution is the distribution that allocates all of its probability to a single outcome. Here, the entropy is zero, since samples from the distribution are always the same – there is no information to be transmitted about a specific sample. On the flip side, the highest entropy distribution is one that is uniform over all outcomes. Here, since all outcomes are

equally likely, even the most efficient code has to convey which of several possibilities was actually chosen at a given sample. Other distributions fall in between, as measured by Eq. 2. In the context of our clustering problem, this maps onto the intuition that it is easy to remember to always put every object in the same cluster (a low entropy distribution, lower representational cost), but harder to remember different clusters for each object – with the extreme being to have a separate cluster for each object (a high entropy distribution, higher representational cost).

We want to use this measure of complexity to inform a probability distribution over mappings π, such that low complexity mappings are preferred over higher complexity ones. If the number of possible clusters is infinite, the entropy can grow arbitrarily large. As in previous work (drawing upon the maximum entropy formulation of the free energy functional; Ortega et al., 2015), we translate this into a probability distribution by taking its weighted negative exponent:

$$P(\pi) = \frac{\exp(-k\mathcal{H}(\pi))}{\sum_{\pi'} \exp(-k\mathcal{H}(\pi'))}$$
$$\propto \exp(-k\mathcal{H}(\pi)) \tag{3}$$

where k is a positive constant that determines how peaked the distribution is and the normalizing factor sums over all possible mappings π'. In the following sections, we show that the CRP corresponds to exactly such a probability distribution.

The relationship between the CRP and entropy

We first discuss the key properties of the CRP. The key property of the CRP is the way a new object is added to an existing clustering of states:

1. Assign it to an existing cluster with probability proportional to the number of items already in the cluster.

2. Assign it to a new cluster with fixed probability α.

This "rich-get-richer" or "preferential attachment" property is inherent in entropy minimization as well. Adding an object to a cluster that already has high probability reduces the entropy of the distribution by making it peakier. On the other hand, adding it to a less populated one moves the distribution closer to uniform, increasing its entropy. Therefore, adding a new object to a cluster that already has many objects in it results in less cost for representing that distribution than adding it to one that has fewer objects.

We can formalize this intuition. By Equation 3, the entropy of a mapping specifies its probability. We can compute the entropies of all the mappings that arise from different possible assignment of a new object to an existing mapping. Inserting these entropies into Equation 3 specifies a probability distribution over the possible assignments of the new object. This way of assigning new objects to clusters is not arbitrary. Rather, it is normative, under the resource-rational assumption that we want to best utilize limited representational resources and thereby prefer mappings with low complexity (and hence low representation cost).

In the Appendix, we provide the mathematical details of the above procedure, for assigning new objects based on probability under Equation 3. We prove that this procedure exactly recovers the CRP's new object assignment rules when the number of objects being classified is reasonably large, for $\alpha = 1/e \approx 0.368$. The constant k (from Equation 3) is given by N, the number of objects. Under this probability distribution, low entropy mappings are more assiduously preferred when the number of objects becomes larger.

Summary

Our goal was to examine the consequences of limited cognitive resources on the clustering process. We find that a prior over clusters proportional to the negative exponent of the entropy, weighted by the number of objects, gives CRP-like clustering. This prior strongly prefers lower entropy mappings to higher entropy ones. This indicates that CRP-like clustering might come from a tendency to reduce representational burden. In

other words, CRP-like clustering can come from a bias toward learning 'simple' object-category mappings, where simplicity is defined as the entropy of the marginal over categories. This provides a process-level theory for why CRP-like priors might be appropriate for modeling human cognition.

Simulations

Our mathematical results establish a direct correspondence between reducing the representational cost of a clustering and the CRP, in the limits of a large number of objects. To determine whether the clustering produced by this resource-rational clustering scheme produces results similar to those expected from a CRP with realistic samples, we conducted a series of simulations where we generated cluster assignments for both distributions and then analyze the correspondence.

The correspondence between the CRP and the entropy-based distribution is closer as the number of objects (N) increases. At very low N therefore, these distributions deviate slightly (see Appendix for details). Since subsequent clustering behavior is conditioned strongly on the object assignments thus far, these differences can amplify. That is, even though the conditional distributions get closer with higher N, the marginal distributions deviate further. The resultant distribution is still qualitatively very similar to the CRP (as discussed below), and it is an interesting direction of future work to examine whether this distribution might *better* describe human clustering behavior than traditional CRPs. Here, however, to validate the similarities with the CRP, we control for this deviation by clustering the first M objects according to an exact CRP.

We evaluated the correspondence between the CRP and our resource-rational distribution based on two criteria. First, a property characteristic of scale-free distributions like the CRP is that the the sizes of the different clusters decay as a power law. Therefore, if we sort the clusters by size, the logarithm of the cluster size (i.e. the fraction of the total number of objects that are in that cluster) should be a linear function of the logarithm of

the cluster index. Second, another key property of the CRP is that the number of clusters increases logarithmically with the total number of objects. We can also measure that for our entropy-based distribution and examine whether the number of clusters is a linear function of the logarithm of the number of objects.

Method

We generated samples from our distribution as follows. We first cluster M objects according to a CRP with $\alpha = 0.368$, varying M between 0 and 80. We then cluster an additional 10^6 objects from this starting point, either with the CRP's conditional cluster assignment rules, or based on the entropy as specified by Eq 3. We cluster such a large number of objects to get a reasonable number of total clusters so we can better analyses the distributions of objects across them – even with this many objects, the average number of total clusters is around 6. This procedure is repeated to get 20 unique set of clusters for each cluster assignment rule and each M.
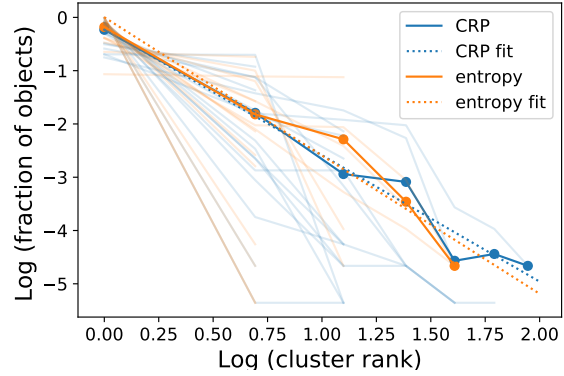
Results

We find that the entropy-based distribution produces clusterings that have statistical properties that closely resemble those of the CRP.

To provide an initial illustration of the correspondence, we focus on the case where $M = 10$. Figure 2A shows how the size of the clusters decays as a power law over the cluster index ranked by size. Figure 2B plots these in log-space and highlights the linear relationship characteristic of a power law. Figure 2C shows the change in the average number of clusters (over the 20 runs) as a function of the logarithm of the number of objects. We see that this grows close to linearly, with the linear fit (dotted line) closely matching the data. The deviation is most apparent at smaller numbers of objects, as expected since the linear relationship is expected in the limit. The number of clusters from the CRP is slightly but not statistically significantly higher than that from the entropy-based clustering.
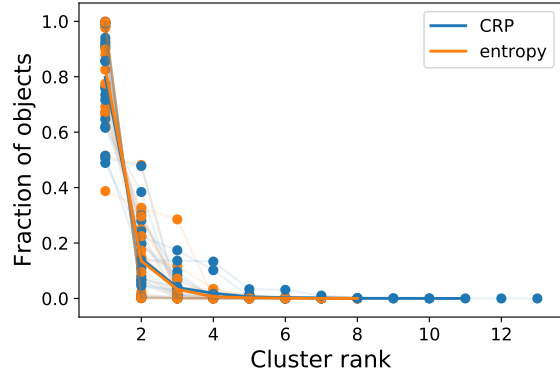
(A) Cluster size as a function of cluster rank.

(B) Log cluster size as a function of log rank.
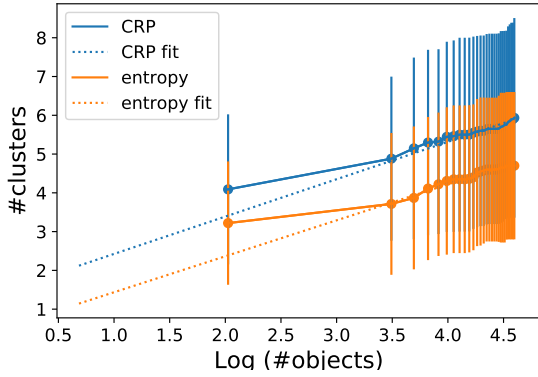
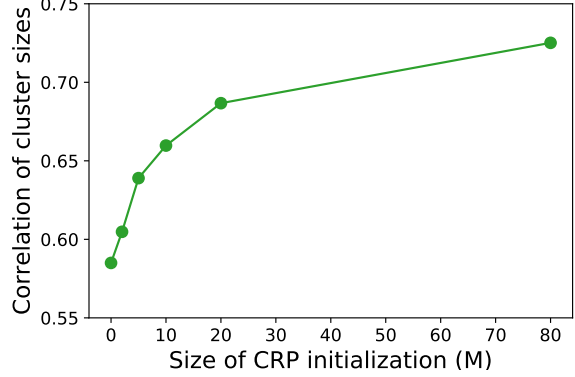(C) Number of clusters vs. log number of objects.

(D) Correlation in cluster sizes with increasing M.

Figure 2. **Simulation results:** A) Cluster size scales as a power law with cluster rank. B) This relationship is highlighted in log space. The average behavior of the two clustering assignments resembles the linear fit (dotted lines). C) The number of clusters grows logarithmically with the number of objects. D) The correlation between the cluster sizes from the two clusterings increases with M, but levels off.

To provide a more detailed picture of this correspondence, we examined how it changes as a function of M. We computed the correlation between the cluster sizes (i.e. the data plotted in Figure 2A, matched by cluster index) between the CRP and the entropy-based clustering for the first 500 objects, after initializing with variable M. We expect the correlation to improve as M increases. We see that this correlation does in fact increase with increasing M (Figure 2D), but the difference is not very large (varies from ~ 0.58 to ~ 0.73) and appears to level off. This indicates that the correspondence between the CRP and the entropy-based prior is fairly robust to the value of the initialization M.

Discussion

Needing to cluster experiences together is a ubiquitous aspect of human cognition. In this paper, we have approached this problem from the perspective of rational analysis, asking how an ideal agent should seek to use their limited cognitive resources. Our results show that the solution to this problem, when those resources are expressed in information-theoretic terms, has a direct correspondence with an approach to clustering that has been widely used in probabilistic models of cognition (the Chinese restaurant process, or CRP). These results provide the first cognitively motivated justification for that assumption.

Our findings suggest interesting directions for future empirical work. If CRP-like clustering comes from representational costs, manipulating these costs should result in different clustering behaviors. Our model predicts that having more limited cognitive resources should affect clustering behavior, driving toward a lower entropy representation and a stronger preference for few, large, clusters. This would not be predicted by a traditional CRP model, since it is a consequence of the cognitive resources available and not a change in the beliefs of the agent about the relative prior probability of different clusterings.

Another broad direction for future research is exploring whether other ubiquitous priors assumed in probabilistic models of human cognition might arise from the algorithmic processes involved in learning and representation. The field of probabilistic numerics (Hennig, Osborne, & Girolami, 2015) describes how numerical processes at the algorithmic level can imitate complex priors at the computational level. This raises the question of the epistemic value of the prior in Bayesian models of cognition – whether it represents pre-existing knowledge, or emergent properties of the algorithm. Our work highlights that this difference can be nuanced.

Acknowledgments

References

Adamic, L. A., & Huberman, B. A. (2002). Zipf's law and the internet. *Glottometrics*, *3*(1), 143–150.

Aldous, D. J. (1985). Exchangeability and related topics. *École d'Été de Probabilités de Saint-Flour XIII—1983*, *1117*, 1–198.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409.

Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509–512.

Bhui, R., & Gershman, S. J. (2018). Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*, *125*(6), 985.

Gershman, S. J., & Blei, D. M. (2012). A tutorial on bayesian nonparametric models. *Journal of Mathematical Psychology*, *56*(1), 1–12.

Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*(1), 197.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Goldwater, S., Griffiths, T. L., & Johnson, M. (2009). A bayesian framework for word segmentation: Exploring the effects of context. *Cognition*, *112*(1), 21–54.

Gottwald, S., & Braun, D. A. (2019). Bounded rational decision-making from elementary computations that reduce uncertainty. *Entropy*, *21*(4), 375.

Griffiths, T. L., & Austerweil, J. (2008). Analyzing human feature learning as nonparametric bayesian inference. *Advances in Neural Information Processing Systems*, *21*, 97–104.

Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in*

Cognitive Sciences, *7*(2), 217–229.

Hennig, P., Osborne, M. A., & Girolami, M. (2015). Probabilistic numerics and
   uncertainty in computations. *Proceedings of the Royal Society A: Mathematical,
   Physical and Engineering Sciences*, *471*(2179), 20150142.

Hjort, N. L., Holmes, C., Müller, P., & Walker, S. G. (2010). *Bayesian nonparametrics*
   (Vol. 28). Cambridge University Press.

Kemp, C., Tenenbaum, J. B., Niyogi, S., & Griffiths, T. L. (2010). A probabilistic model of
   theory formation. *Cognition*, *114*(2), 165–196.

Laughlin, S. (1981). A simple coding procedure enhances a neuron's information capacity.
   Zeitschrift für Naturforschung c, *36*(9-10), 910–912.

Lieder, F., & Griffiths, T. L. (2019). Resource-rational analysis: Understanding human
   cognition as the optimal use of limited computational resources. *Behavioral and
   Brain Sciences*, 1–85.

Mandelbrot, B. (1960). The pareto-levy law and the distribution of income. *International
   Economic Review*, *1*(2), 79–106.

Olshausen, B. A., & Field, D. J. (1996). Natural image statistics and efficient coding.
   Network: Computation in Neural Systems, *7*(2), 333–339.

Ortega, P. A., Braun, D. A., Dyer, J., Kim, K. E., & Tishby, N. (2015).
   Information-theoretic bounded rationality. *arXiv preprint arXiv:1512.06789*.

Rosen, K. T., & Resnick, M. (1980). The size distribution of cities: an examination of the
   pareto law and primacy. *Journal of Urban Economics*, *8*(2), 165–186.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell Labs Technical
   Journal*, *27*(3), 379–423.

Todorov, E. (2009). Efficient computation of optimal actions. *Proceedings of the National
   Academy of Sciences*, *106*(28), 11478–11483.

Zenon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on
   the costs of cognition. *Neuropsychologia*, *123*, 5–18.

Appendix

Formal definition of the CRP

Given a cluster assignment π, an object is added to give π'. This can be split into two cases, where in the object is added to cluster j to give π_j or if it is added to a new cluster (i.e. one with $n_j = 0$) to give π^0. Formalizing the CRP in these terms:

$$P(\pi_j|\pi) \propto n_j$$
$$P(\pi^0|\pi) \propto \alpha \tag{4}$$

How does new item assignment affect the entropy of the mapping? We start with a given π, which specifies a cluster assignment of N objects into N clusters, with each cluster containing n_i objects. The resulting entropy is:

$$\begin{aligned}
\mathcal{H} &= -\sum_i^N \frac{n_i}{N} \log\frac{n_i}{N} \\
&= -\sum_i^K \frac{n_i}{N} \log n_i + \log N
\end{aligned} \tag{5}$$

We then consider how the entropy changes when the new item is assigned to give a new π'.

Adding to an existing cluster

When adding a new object, we can add an object to an existing cluster j to give $\mathcal{H}_j$:

$$\mathcal{H}_j = -\sum_{i\backslash j}^K \frac{n_i}{N+1} \log n_i - \frac{n_j+1}{N+1} \log(n_j+1) + \log(N+1)$$

The difference in the entropy between the new and the old distributions is:

$$\mathcal{H}_j - \mathcal{H} = \sum_{i \backslash j}^{K} n_i \left(\frac{1}{N} - \frac{1}{N+1} \right) \log n_i$$
$$+ \frac{n_j}{N} \log n_j - \frac{n_j + 1}{N+1} \log(n_j + 1)$$
$$+ \log(N+1) - \log N$$

We want to separate out the terms dependent on n_j, so we separate out the first term as:

$$\sum_{i \backslash j}^{K} n_i \left(\frac{1}{N} - \frac{1}{N+1} \right) \log n_i$$
$$= \sum_{i}^{K} \frac{n_i \log n_i}{N(N+1)}$$
$$- \frac{n_j \log n_j}{N(N+1)}$$

The difference therefore reduces to the following terms, with E representing the terms independent of n_j:

$$\mathcal{H}_j - \mathcal{H} = E + \frac{n_j}{N} \log n_j - \frac{n_j + 1}{N+1} \log(n_j + 1) - \frac{n_j}{N(N+1)} \log n_j$$

Here, the independent of n_j term denoted E is given by:

$$E = \sum_{i}^{K} \frac{n_i \log n_i}{N(N+1)} + \log(N+1) - \log N$$

We further simplify it:

$$\frac{n_j}{N} \log n_j - \frac{n_j + 1}{N + 1} \log(n_j + 1) - \frac{n_j}{N(N + 1)} \log n_j$$

$$= \frac{N n_j \log n_j + n_j \log n_j - N n_j \log(n_j + 1) - N \log(n_j + 1) - n_j \log n_j}{N(N + 1)}$$

$$= -\frac{1}{N(N + 1)} \left[N n_j \log \left(1 + \frac{1}{n_j} \right) - N \log(n_j + 1) \right]$$

$$= -\frac{n_j \log \left(1 + \frac{1}{n_j} \right)}{N + 1} - \frac{\log(n_j + 1)}{N + 1}$$

To get a sense for how this scales, we take the large N limit. In this limit, we only need to consider the leading order terms. In order to do this, we need to make an assumption about the relation between the average number of objects and the total number of objects N. We assume that the average number of objects in a cluster grows grows sub-linearly with the total number of objects. This is a weak assumption and is true as long as a) not all objects are assigned to the same clusters or b) not all objects are assigned to a new cluster. Therefore, both n_j and N grow when N is large, but n_j grows slower. We can therefore Taylor expand the first term as follows:

$$-\frac{n_j \log \left(1 + \frac{1}{n_j} \right)}{N + 1} = -\frac{n_j \left(\frac{1}{n_j} - \frac{1}{2n_j^2} \right)}{N + 1} = \frac{-1}{N + 1} + \frac{1}{2n_j(N + 1)}$$

At large N (and correspondingly large n_j), the leading term (ignore '+ 1's) gives :

$$\mathcal{H}_j - \mathcal{H} \sim E - \frac{\log(n_j)}{N} - \frac{1}{N} \tag{6}$$

Creating a new cluster

If we instead create a new cluster, the entropy gives:

$$\mathcal{H}^0 = -\sum_i^K \frac{n_i}{N+1} \log n_i - \frac{1}{N+1} \log 1 + \log(N+1)$$

$$= -\sum_i^K \frac{n_i}{N+1} \log n_i + \log(N+1)$$

Computing the entropy difference, we get the same expression except with $n_j = 0$:

$$\mathcal{H}^0 - \mathcal{H} = E$$

From entropy to probability distribution

As discussed on the main text, we consider the negative exponential functional form for the prior distribution:

$$P(\pi) \propto \exp(-N\mathcal{H}(\pi))$$

We show that this functional form satisfies the CRP assignment rules in Equation 4. Rewriting the conditional distributions in Equation 4 as ratios, we get:

$$\frac{P(\pi_j)}{P(\pi)} \propto n_j$$

$$\frac{P(\pi^0)}{P(\pi)} \propto \alpha$$

For the first constraint:

$$\frac{P(\pi_j)}{P(\pi)} \propto \exp\left(-N(\mathcal{H}_j - \mathcal{H})\right)$$

$$= \exp\left[-N\left(E - \frac{\log(n_j)}{N} - \frac{1}{N}\right)\right]$$

$$= \exp(-NE) \times n_j \times e$$

$$\propto n_j \tag{7}$$

To verify the second constraint,

$$\frac{P(\pi^0)}{P(\pi)} \propto \exp\left(-N(\mathcal{H}^0 - \mathcal{H})\right)$$

$$= \exp\left(-NE\right)$$

We can then normalize the probability of the new clusterings as follows:

$$P(\pi_j) = P(\pi)\frac{n_j}{N + 1/e}$$

$$P(\pi^0) = P(\pi)\frac{1/e}{N + 1/e}$$

This is equivalent to a CRP with $\alpha = 1/e \approx 0.36$. Note that we can get a corresponding CRP with a different α by taking the logarithm in the Equation 2 and the exponent in Equation 7 with respect to a different constant.

We also revisit the approximation made in Equation 6. If we had not made the approximation required to eliminate the extra term, we would have an additional dependence on n_j as follows:

$$\frac{P(\pi_j)}{P(\pi)} \propto \exp\left(-N(\mathcal{H}_j - \mathcal{H})\right)$$

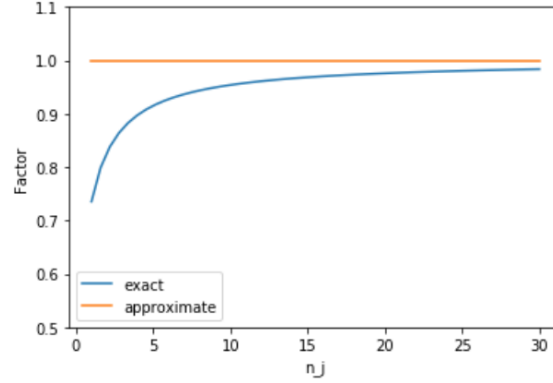$$= \exp(-NE) \times n_j \times \exp(n_j \log(1 + 1/n_j))$$

Figure A1. **Evaluating the approximation**: We plot the terms in the exponents of Equation 8 as a function of increasing number of objects in a cluster (n_j). We see that the difference between the approximation and exact value reduces quickly.

In our simplification, we are making the following approximation:

$$\exp(n_j\log(1 + 1/n_j)) \sim \exp(1) \tag{8}$$

We plot these exponents in Figure A1 to give a sense for when this is a good approximation. We see that even at small n_j, the values are relatively close, with the approximation converging quickly.

We restrict our analyses to the correspondence of the conditional distributions $P(\pi_N|\pi_{N-1})$ between the entropy-based distribution and the CRP, rather than directly examining the joint distribution $P(\pi_N)$. This is because computing the normalization factor for the conditional distribution for the entropy-based distribution (*before* making the approximation above) depends on the distribution of objects in the previous step – unlike after we make the approximation when the normalization factor goes to $Ne + 1$. This makes the pre-approximation joint distribution difficult to compute.