

Optimal Number of Choices in Rating Contexts

Sam Ganzfried¹ and Farzana Yusuf²

¹ Ganzfried Research

² School of Computing and Information Sciences, Florida International University
sam@ganzfriedresearch.com, fyusu003@fiu.edu

Abstract

In many settings people must give numerical scores to entities from a small discrete set. For instance, rating physical attractiveness from 1–5 on dating sites, or papers from 1–10 for conference reviewing. We study the problem of understanding when using a different number of options is optimal. We consider the case when scores are uniform random and Gaussian. We study when using 2, 3, 4, 5, and 10 options out of a total of 100 is optimal in these models. One may expect that using more options would always improve performance in this model, but we show that this is not necessarily the case, and that using fewer choices—even just two—can surprisingly be optimal in certain situations. While in theory for this setting it would be optimal to use all 100 options, in practice this is prohibitive, and it is preferable to utilize a smaller number of options due to humans’ limited computational resources. Our results could have many potential applications, as settings requiring entities to be ranked by humans are ubiquitous. There could also be applications to other fields such as signal or image processing where input values from a large set must be mapped to output values in a smaller set.

Introduction

Humans rate items or entities in many important settings. For example, users of dating websites and mobile applications rate other users’ physical attractiveness, teachers rate scholarly work of students, and reviewers rate the quality of academic conference submissions. In these settings, the users assign a numerical (integral) score to each item from a small discrete set. However, the number of options in this set can vary significantly between applications, and even within different instantiations of the same application. For instance, for rating attractiveness, three popular sites all use a different number of options. On “Hot or Not,” users rate the attractiveness of photographs submitted voluntarily by other users on a scale of 1–10 (Figure 1¹). These scores are aggregated and the average is assigned as the overall “score” for a photograph. On the dating website OkCupid, users rate other users on a scale of 1–5 (if a user rates another user 4

or 5 then the rated user receives a notification)² (Figure 2³). And on the mobile application Tinder users “swipe right” (green heart) or “swipe left” (red X) to express interest in other users (two users are allowed to message each other if they mutually swipe right), which is essentially equivalent to using a binary $\{1, 2\}$ scale (Figure 3⁴). Education is another important application area requiring human ratings. For the 2016 International Joint Conference on Artificial Intelligence, reviewers assigned a “Summary Rating” score from –5–5 (equivalent to 1–10) for each submitted paper (Figure 4)⁵. The papers are then discussed and scores aggregated to produce an acceptance or rejection decision based on the average of the scores.

Despite the importance and ubiquity of the problem, there has been little fundamental research done on the problem of determining the optimal number of options to allow in such settings. We study a model in which users have a underlying integral ground truth score for each item in $\{1, \dots, n\}$ and are required to submit an integral rating in $\{1, \dots, k\}$, for $k \ll n$. (For ease of presentation we use the equivalent formulation $\{0, \dots, n-1\}$, $\{0, \dots, k-1\}$.) We use two generative models for the ground truth scores: a uniform random model in which the fraction of scores for each value from 0 to $n-1$ is chosen uniformly at random (by choosing a random value for each and then normalizing), and a model where scores are chosen according to a Gaussian distribution with a given mean and variance. We then compute a “compressed” score distribution by mapping each full score s from $\{0, \dots, n-1\}$ to $\{0, \dots, k-1\}$ by applying

$$s \leftarrow \left\lfloor \frac{s}{\binom{n}{k}} \right\rfloor. \quad (1)$$

We then compute the average “compressed” score a_k , and

²The likelihood of receiving an initial message is actually much more highly correlated with the variance—and particularly the number of “5” ratings—than with the average rating (Fry 2015).

³<http://blog.okcupid.com/index.php/the-mathematics-of-beauty/>

⁴<https://tctechcrunch2011.files.wordpress.com/2015/11/tinder-two.jpg>

⁵<https://easychair.org/conferences/?conf=ijcail6>

¹<http://blog.mrmeyer.com/2007/are-you-hot-or-not/>



Figure 1: Hot or Not users rate attractiveness 1–10.




Figure 2: OkCupid users rate attractiveness 1–5.



Figure 3: Tinder users rate attractiveness 1–2.

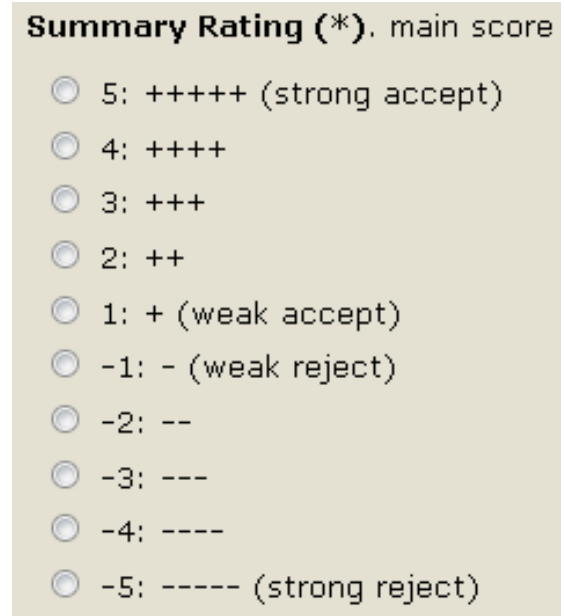


Figure 4: IJCAI reviewers rate papers -5–5.

compute its error e_k according to

$$e_k = \left| a_f - \frac{n-1}{k-1} \cdot a_k \right|, \quad (2)$$

where a_f is the ground truth average. The goal is to pick $\arg\min_k e_k$. While there are many possible generative models and cost functions, these seemed like the most natural.

We derive a closed-form expression for e_k that depends on only a small number (k) of parameters of the underlying distribution for an arbitrary distribution.⁶ This allows us

⁶For theoretical simplicity we theoretically study a continuous version where scores are chosen according to a distribution over $(0, n)$ (though the simulations are for the discrete version) and the compressed scores are over $\{0, \dots, k-1\}$. In this setting we use a normalization factor of $\frac{n}{k}$ instead of $\frac{n-1}{k-1}$ for the e_k term. Continuous approximations for large discrete spaces have been studied in other settings; for instance, they have led to simplified analysis

to exactly characterize the performance of using each number of choices. In computational simulations we repeatedly compute e_k and compare the average values. We focus on $n = 100$ and $k = 2, 3, 4, 5, 10$, which we believe are the most natural and interesting choices for initial study.

One could argue that this model is somewhat “trivial” in the sense that it would be optimal to set $k = n$ to permit all the possible scores, as this would result in the “compressed” scores agreeing exactly with the full scores. However, there are several reasons that would lead us to prefer to select $k \ll n$ in practice (as all of the examples previously described have done), thus making this analysis worthwhile. It is much easier for a human to assign a score from a small set than from a large set, particularly when rating many items under time constraints. We could have included an additional

and insight in poker games with continuous distributions of private information (Ankenman and Chen 2006).

term into the cost function e_k that explicitly penalizes larger values of k , which would have a significant effect on the optimal value of k (providing a favoritism for smaller values). However the selection of this function would be somewhat arbitrary and would make the model more complex, and we leave this for future study. Given that we do not include such a penalty term, one may expect that increasing k will always decrease e_k in our setting. While the simulations show a clear negative relationship, we show that smaller values of k actually lead to smaller e_k surprisingly often. These smaller values would receive further preference with a penalty term.

The most closely related theoretical work studies the impact of using finely grained numerical grades (100, 99, 98) vs. coarse letter grades (A, B, C) (Dubey and Geanakoplos 2010). They conclude that if students care primarily about their rank relative to the other students, they are often best motivated to work by assigning them coarse categories than exact numerical scores. In a setting of “disparate” student abilities they show that the optimal absolute grading scheme is always coarse. Their model is game-theoretic; each player (student) selects an effort level, seeking to optimize a utility function that depends on both the relative score and effort level. Their setting is quite different from ours in many ways. For one, they study a setting where it is assumed that the underlying “ground truth” score is known, yet may be disguised for strategic reasons. In our setting the goal is to approximate the ground truth score as closely as possible.

While we are not aware of prior theoretical study of our exact problem, there have been experimental studies on the optimal number of options on a “Likert scale” (Likert 1932; Matell and Jacoby 1971; Wildt and Mazis 1978; Cox III 1980; Friedman, Wilamowsky, and Friedman 1981). The general conclusion is that “the optimal number of scale categories is content specific and a function of the conditions of measurement.” (Garland 1991) There has been study of whether including a “mid-point” option (i.e., the middle choice from an odd number) is beneficial. One experiment demonstrated that the use of the mid-point category decreases as the number of choices increases: 20% of respondents choose the mid-point for 3 and 5 options while only 7% did for 7, 9, . . . , 19 (Matell and Jacoby 1972). They conclude that it is preferable to either not include a mid-point at all or use a large number of options. Subsequent experiments demonstrated that eliminating a mid-point can reduce social desirability bias which results from respondents’ desires to please the interviewer or not give a perceived socially unacceptable answer (Garland 1991). There has also been significant research on questionnaire design and the concept of “feeling thermometers,” particularly from the fields of psychology and sociology (Sudman, Bradburn, and Schwartz 1996; McKennell 1974; Lewis and Erdinç 2017; C Preston and Colman 2000; Lozano, Garca-Cueto, and Muiz 2008; R. Lehmann and Hulbert 1972). One study concludes from experimental data: “in the measurement of satisfaction with various domains of life, 11-point scales clearly are more reliable than comparable 7-point scales” (Alwin 1997). Another study shows that “people are more likely to purchase gourmet jams or chocolates or to undertake optional class essay assignments when offered a limited array

of 6 choices rather than a more extensive array of 24 or 30 choices” (S. Iyengar and Lepper 2001). Since the experimental conclusions are dependent on the specific datasets and seem to vary from domain to domain, we choose to focus on formulating theoretical models and computational simulations, though we also include results and discussion from several datasets in the appendix.

Theoretical characterization

Suppose scores are given by continuous pdf f (with cdf F) on $(0, 100)$, and we wish to compress them to two options, $\{0, 1\}$. Scores below 50 are mapped to 0, and above 50 to 1. The average of the full distribution is $a_f = E[X] = \int_{x=0}^{100} x f(x) dx$. The average of the compressed version is

$$a_2 = \int_{x=0}^{50} 0 f(x) dx + \int_{x=50}^{100} 1 f(x) dx = 1 - F(50).$$

So $e_2 = |a_f - 100(1 - F(50))| = |E[X] - 100 + 100F(50)|$. For three options,

$$\begin{aligned} a_3 &= \int_{x=0}^{\frac{100}{3}} 0 f(x) dx + \int_{x=\frac{100}{3}}^{\frac{200}{3}} 1 f(x) dx + \int_{x=\frac{200}{3}}^{100} 2 f(x) dx \\ &= 2 - F(100/3) - F(200/3) \\ e_3 &= |a_f - 50(2 - F(100/3) - F(200/3))| \\ &= |E[X] - 100 + 50F(100/3) + 50F(200/3)| \end{aligned}$$

In general for n total and k compressed options,

$$\begin{aligned} a_k &= \sum_{i=0}^{k-1} \int_{x=\frac{ni}{k}}^{\frac{n(i+1)}{k}} i f(x) dx \\ &= \sum_{i=0}^{k-1} \left[i \left(F\left(\frac{n(i+1)}{k}\right) - F\left(\frac{ni}{k}\right) \right) \right] \\ &= (k-1)F(n) - \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \\ &= (k-1) - \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \\ e_k &= \left| a_f - \frac{n}{k-1} \left((k-1) - \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \right) \right| \\ &= \left| E[X] - n + \frac{n}{k-1} \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \right| \quad (3) \end{aligned}$$

Equation 3 allows us to characterize the relative performance of choices of k for a given distribution f . For each k it requires only knowing k statistics of f (the $k-1$ values of $F\left(\frac{ni}{k}\right)$ plus $E[X]$). In practice these could likely be closely approximated from historical data for small k values.

As an example we see that $e_2 < e_3$ iff

$$|E[X] - 100 + 100F(50)| < \left| E[X] - 100 + 50F\left(\frac{100}{3}\right) + 50F\left(\frac{200}{3}\right) \right|$$

Consider a full distribution that has half its mass right around 30 and half its mass right around 60 (Figure 5). Then $a_f = E[X] = 0.5 \cdot 30 + 0.5 \cdot 60 = 45$. If we use $k = 2$, then the mass at 30 will be mapped down to 0 (since $30 < 50$) and the mass at 60 will be mapped up to 1 (since $60 > 50$) (Figure 6). So $a_2 = 0.5 \cdot 0 + 0.5 \cdot 1 = 0.5$. Using normalization of $\frac{n}{k} = 100$, $e_2 = |45 - 100(0.5)| = |45 - 50| = 5$. If we use $k = 3$, then the mass at 30 will also be mapped down to 0 (since $0 < \frac{100}{3}$); but the mass at 60 will be mapped to 1 (not the maximum possible value of 2 in this case), since $\frac{100}{3} < 60 < \frac{200}{3}$ (Figure 6). So again $a_3 = 0.5 \cdot 0 + 0.5 \cdot 1 = 0.5$, but now using normalization of $\frac{n}{k} = 50$ we have $e_3 = |45 - 50(0.5)| = |45 - 25| = 20$. So, surprisingly, in this example allowing more ranking choices actually significantly increases error.




Figure 5: Example distribution for which compressing with $k = 2$ produces lower error than $k = 3$.

If we happened to be in the case where both $a_2 \leq a_f$ and $a_3 \leq a_f$, then we could remove the absolute values and reduce the expression to see that $e_2 < e_3$ iff $\int_{x=\frac{100}{3}}^{50} f(x)dx < \int_{x=\frac{200}{3}}^{50} f(x)dx$. One could perform more comprehensive analysis considering all cases to obtain better characterization and intuition for the optimal value of k for distributions with different properties.

Rounding compression

An alternative model we could have considered is to use rounding to produce the compressed scores as opposed to using the floor function from Equation 1. For instance, for the case $n = 100, k = 2$, instead of dividing s by 50 and taking the floor, we could instead partition the points according to whether they are closest to $t_1 = 25$ or $t_2 = 75$. In the example above, the mass at 30 would be mapped to t_1 and the mass at 60 would be mapped to t_2 . This would produce a compressed average score of $a_2 = \frac{1}{2} \cdot 25 + \frac{1}{2} \cdot 75 = 50$. No normalization would be necessary, and this would produce error of $e_2 = |a_f - a_2| = |45 - 50| = 5$, as the floor approach did as well. Similarly, for $k = 3$ the region midpoints will be $q_1 = \frac{100}{6}, q_2 = 50, q_3 = \frac{500}{6}$. The mass at 30 will be mapped to $q_1 = \frac{100}{6}$ and the mass at 60 will be mapped to $q_2 = 50$. This produces a compressed average score of

$a_3 = \frac{1}{2} \cdot \frac{100}{6} + \frac{1}{2} \cdot 50 = \frac{100}{3}$. This produces an error of $e_3 = |a_f - a_3| = |45 - \frac{100}{3}| = \frac{35}{3} = 11.67$. Although the error for $k = 3$ is smaller than for the floor case, it is still significantly larger than $k = 2$'s, and using two options still outperforms using three for the example in this new model.

In general, this approach would create k "midpoints" $\{m_i^k\}$: $m_i^k = \frac{n(2i-1)}{2k}$. For $k = 2$ we have

$$\begin{aligned} a_2 &= \int_{x=0}^{50} 25 + \int_{x=50}^{100} 75 = 75 - 50F(50) \\ e_2 &= |a_f - (75 - 50F(50))| = |E[X] - 75 + 50F(50)| \end{aligned}$$

One might wonder whether the floor approach would ever outperform the rounding approach (in the example above the rounding approach produced lower error $k = 3$ and the same error for $k = 2$). As a simple example to see that it can, consider the distribution with all mass on 0. The floor approach would produce $a_2 = 0$ giving an error of 0, while the rounding approach would produce $a_2 = 25$ giving an error of 25. Thus, the superiority of the approach is dependent on the distribution. We explore this further in the experiments.

For three options,

$$\begin{aligned} a_3 &= \int_0^{\frac{100}{3}} \frac{100}{6} f(x) + \int_{\frac{100}{3}}^{\frac{200}{3}} 50 f(x) + \int_{\frac{200}{3}}^{100} \frac{500}{6} f(x) \\ &= \frac{500}{6} - \frac{100}{3} F\left(\frac{100}{3}\right) - \frac{100}{3} F\left(\frac{200}{3}\right) \\ e_3 &= \left| E[X] - \frac{500}{6} + \frac{100}{3} F\left(\frac{100}{3}\right) + \frac{100}{3} F\left(\frac{200}{3}\right) \right| \end{aligned}$$

For general n and k , analysis as above yields

$$\begin{aligned} a_k &= \sum_{i=0}^{k-1} \int_{x=\frac{ni}{k}}^{\frac{n(i+1)}{k}} m_{i+1}^k f(x) dx = \frac{n(2k-1)}{2k} - \frac{n}{k} \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \\ e_k &= \left| a_f - \left[\frac{n(2k-1)}{2k} - \frac{n}{k} \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \right] \right| \quad (4) \\ &= \left| E[X] - \frac{n(2k-1)}{2k} + \frac{n}{k} \sum_{i=1}^{k-1} F\left(\frac{ni}{k}\right) \right| \quad (5) \end{aligned}$$

Like for the floor model e_k requires only knowing k statistics of f . The rounding model has an advantage over the floor model that there is no need to convert scores between different scales and perform normalization. One drawback is that it requires knowing n (the expression for m_i^k is dependent on n), while the floor model does not. In our experiments we assume $n = 100$, but in practice it may not be clear what the agents' ground truth granularity is and may be easier to just deal with scores from 1 to k . Furthermore, it may seem unnatural to essentially ask people to rate items as " $\frac{100}{6}, 50, \frac{500}{6}$ ", rather than "1, 2, 3" (though the conversion between the score and m_i^k could be done behind the scenes essentially circumventing the potential practical complication). One can generalize both the floor and rounding model by using a score of $s(n, k)_i$ for the i 'th region. For the floor setting we set $s(n, k)_i = i$, and for the rounding setting $s(n, k)_i = m_i^k = \frac{n(2i+1)}{2k}$.





Figure 6: Compressed distributions using $k = 2$ and $k = 3$ for example from Figure 5.

Computational simulations

The above analysis leads to the immediate question of whether the example for which $e_2 < e_3$ was a fluke or whether using fewer choices can actually reduce error under reasonable assumptions on the generative model. We study this question using simulations with what we believe are the two most natural models. While we have studied the continuous setting where the full set of options is continuous over $(0, n)$ and the compressed set is discrete $\{0, \dots, k-1\}$, we now consider the perhaps more realistic setting where the full set is the discrete set $\{0, \dots, n-1\}$ and the compressed set is the same (though it should be noted that the two settings are likely quite similar qualitatively).

The first generative model we consider is a uniform model in which the values of the pmf for each of the n possible values are chosen independently and uniformly at random. The second is a Gaussian model in which the values are generated according to a normal distribution with specified mean μ and standard deviation σ (values below 0 are set to 0 and above $n-1$ to $n-1$). This model also takes as a parameter a number of samples s to use for generating the scores.

For our simulations we used $n = 100$, and considered $k = 2, 3, 4, 5, 10$, which are popular and natural values. For the Gaussian model we used $s = 1000$, $\mu = 50$, $\sigma = \frac{50}{3}$. For each set of simulations we computed the errors for all considered values of k for $m = 100,000$ “items” (each corresponding to a different distribution generated according to the specified model). The main quantities we are interested in computing are the number of times that each value of k produces the lowest error over the m items, and the average value of the errors over all items for each k value.

In the first set of experiments, we compared performance between using $k = 2, 3, 4, 5, 10$ to see for how many of the trials each value of k produced the minimal error (Table 1). Not surprisingly, we see that the number of victories increases monotonically with the value of k , while the average error decreased monotonically (recall that we would have zero error if we set $k = 100$). However, what is perhaps surprising is that using a smaller number of compressed scores produced the optimal error in a far from negligible number of the trials. For the uniform model, using 10 scores

minimized error only around 53% of the time, while using 5 scores minimized error 17% of the time, and even using 2 scores minimized it 5.6% of the time. The results were similar for the Gaussian model, though a bit more in favor of larger values of k , which is what we would expect because the Gaussian model is less likely to generate “fluke” distributions that could favor the smaller values.

	2	3	4	5	10
Uniform # victories	5564	9265	14870	16974	53327
Uniform average error	1.32	0.86	0.53	0.41	0.19
Gaussian # victories	3025	7336	14435	17800	57404
Gaussian average error	1.14	0.59	0.30	0.22	0.10

Table 1: Number of times each value of k in $\{2, 3, 4, 5, 10\}$ produces minimal error and average error values, over 100,000 items generated according to both models.

We next explored the number of victories between just $k = 2$ and $k = 3$, with results in Table 2. Again we observed that using a larger value of k generally reduces error, as expected. However, we find it extremely surprising that using $k = 2$ produces a lower error 37% of the time. As before, the larger k value performs relatively better in the Gaussian model. We also looked at results for the most extreme comparison, $k = 2$ vs. $k = 10$ (Table 3). Using 2 scores outperformed 10 8.3% of the time in the uniform setting, which was larger than we expected. In Figures 7–8, we present a distribution for which $k = 2$ particularly outperformed $k = 10$. The full distribution has mean 54.188, while the $k = 2$ compression has mean 0.548 (54.253 after normalization) and $k = 10$ has mean 5.009 (55.009 after normalization). The normalized errors between the means were 0.906 for $k = 10$ and 0.048 for $k = 2$, yielding a difference of 0.859 in favor of $k = 2$.

We next repeated the extreme $k = 2$ vs. 10 comparison, but we imposed a restriction that the $k = 10$ option could not give a score below 3 or above 6 (Table 4). (If it selected a score below 3 then we set it to 3, and if above 6 we set it to 6). For some settings, for instance paper reviewing, extreme scores are very uncommon, and we strongly suspect that the vast majority of scores are in this middle range. Some possi-

	2	3
Uniform number of victories	36805	63195
Uniform average error	1.31	0.86
Gaussian number of victories	30454	69546
Gaussian average error	1.13	0.58

Table 2: Results for $k = 2$ vs. 3.

	2	10
Uniform number of victories	8253	91747
Uniform average error	1.32	0.19
Gaussian number of victories	4369	95631
Gaussian average error	1.13	0.10

Table 3: Results for $k = 2$ vs. 10.

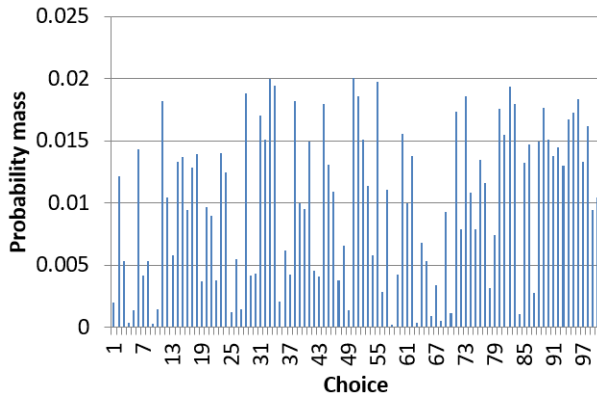


Figure 7: Example distribution where compressing with $k = 2$ produces significantly lower error than $k = 10$. The full distribution has mean 54.188, while the $k = 2$ compression has mean 0.548 (54.253 after normalization) and the $k = 10$ compression has mean 5.009 (55.009 after normalization). The normalized errors between the means were 0.906 for $k = 10$ and 0.048 for $k = 2$, yielding a difference of 0.859 in favor of $k = 2$.

ble explanations are that reviewers who give extreme scores may be required to put in additional work to justify their scores and are more likely to be involved in arguments with other reviewers (or with the authors in the rebuttal). Reviewers could also experience higher regret or embarrassment for being “wrong” and possibly off-base in the review by missing an important nuance. In this setting using $k = 2$ outperforms $k = 10$ nearly $\frac{1}{3}$ of the time in the uniform model.

We also considered the situation where we restricted the $k = 10$ scores to fall between 3 and 7 (as opposed to 3 and 6). Note that the possible scores range from 0–9, so this restriction is asymmetric in that the lowest three possible scores are eliminated while only the highest two are. This is motivated by the intuition that raters may be less inclined to give extremely low scores which may hurt the feelings of an author (for the case of paper reviewing). In this setting, which is seemingly quite similar to the 3–6 setting, $k = 2$ produced lower error 93% of the time in the uniform model!

We next repeated these experiments for the rounding com-

	2	10
Uniform number of victories	32250	67750
Uniform average error	1.31	0.74
Gaussian number of victories	10859	89141
Gaussian average error	1.13	0.20

Table 4: For $k = 10$ only scores 3–6 permitted.

	2	10
Uniform number of victories	93226	6774
Uniform average error	1.31	0.74
Gaussian number of victories	54459	45541
Gaussian average error	1.13	1.09

Table 5: For $k = 10$ only scores 3–7 permitted.

pression function. There are several interesting observations from Table 6. In this setting, $k = 3$ is the clear choice, performing best in both models (by a large margin for the Gaussian model). The smaller values of k perform significantly better with rounding than flooring (as indicated by lower errors) while the larger values perform significantly worse, and their errors seem to approach 0.5 for both models. Taking both compressions into account, the optimal overall approach would still be to use flooring with $k = 10$, which produced the smallest average errors of 0.19 and 0.1 in the two models, while using $k = 3$ with rounding produced errors of 0.47 and 0.24. The 2 vs. 3 experiments produced very similar results for the two compressions (Table 7). The 2 vs. 10 results were quite different, with 2 performing better almost 40% of the time with rounding, vs. less than 10% with flooring (Table 8). In the 2 vs. 10 truncated 3–6 experiments 2 performed relatively better with rounding for both models (Table 9), and for the 2 vs. 10 truncated 3–7 experiments $k = 2$ performed better nearly all the time (Table 10).

	2	3	4	5	10
Uniform # victories	15766	33175	21386	19995	9678
Uniform average error	0.78	0.47	0.55	0.52	0.50
Gaussian # victories	13262	64870	10331	9689	1848
Gaussian average error	0.67	0.24	0.50	0.50	0.50

Table 6: Number of times each value of k produces minimal error and average error values, over 100,000 items generated according to both models with rounding compression.

	2	3
Uniform number of victories	33585	66415
Uniform average error	0.78	0.47
Gaussian number of victories	18307	81693
Gaussian average error	0.67	0.24

Table 7: $k = 2$ vs. 3 with rounding compression.

Conclusion

Settings in which humans must rate items or entities from a small discrete set of options are ubiquitous. We have singled





Figure 8: Compressed distribution for $k = 2$ vs. 10 for example from Figure 7.

	2	10
Uniform number of victories	37225	62775
Uniform average error	0.78	0.50
Gaussian number of victories	37897	62103
Gaussian average error	0.67	0.50

Table 8: $k = 2$ vs. 10 with rounding compression.

	2	10
Uniform number of victories	55676	44324
Uniform average error	0.79	0.89
Gaussian number of victories	24128	75872
Gaussian average error	0.67	0.34

Table 9: $k = 2$ vs. 10 with rounding compression. For $k = 10$ only scores permitted between 3 and 6.

	2	10
Uniform number of victories	99586	414
Uniform average error	0.78	3.50
Gaussian number of victories	95692	4308
Gaussian average error	0.67	1.45

Table 10: $k = 2$ vs. 10 with rounding compression. For $k = 10$ only scores permitted between 3 and 7.

out several important applications—rating attractiveness for dating websites, assigning grades to students, and reviewing academic papers. The number of available options can vary considerably, even within different instantiations of the same application. For instance, we saw that three popular sites for attractiveness rating use completely different systems: Hot or Not uses a 1–10 system, OkCupid uses 1–5 “star” system, and Tinder uses a binary 1–2 “swipe” system.

Despite the problem’s importance, we have not seen it studied theoretically previously. Our goal is to select k to minimize the average (normalized) error between the compressed average score and the ground truth average. We studied two natural models for generating the scores. The first is a uniform model where the scores are selected independently and uniformly at random, and the second is a Gaus-

sian model where they are selected according to a more structured procedure that gives preference for the options near the center. We provided a closed-form solution for continuous distributions with arbitrary cdf. This allows us to characterize the relative performance of choices of k for a given distribution. We saw that, counterintuitively, using a smaller value of k can actually produce lower error for some distributions (even though we know that as k approaches n the error approaches 0): we presented specific distributions for which using $k = 2$ outperforms 3 and 10.

We performed numerous simulations comparing the performance between different values of k for different generative models and metrics. The main metric was the absolute number of times for which values of k produced the minimal error. We also considered the average error over all simulated items. Not surprisingly, we observed that performance generally improves monotonically with k as expected, and more so for the Gaussian model than uniform. However, we observe that small k values can be optimal a non-negligible amount of the time, which is perhaps counterintuitive. In fact, using $k = 2$ outperformed $k = 3, 4, 5$, and 10 on 5.6% of the trials in the uniform setting. Just comparing 2 vs. 3, $k = 2$ performed better around 37% of the time. Using $k = 2$ outperformed 10 8.3% of the time, and when we restricted $k = 10$ to only assign values between 3 and 7 inclusive, $k = 2$ actually produced lower error 93% of the time! This could correspond to a setting where raters are ashamed to assign extreme scores (particularly extreme low scores). We compared two different natural compression rules—one based on the floor function and one based on rounding—and weighed the pros and cons of each. For smaller k rounding leads to significantly lower error than flooring, with $k = 3$ the clear optimal choice, while for larger k rounding leads to much larger error.

A future avenue is to extend our analysis to better understand specific distributions for which different k values are optimal, while our simulations are in aggregate over many distributions. Application domains will have distributions with different properties, and improved understanding will allow us to determine which k is optimal for the types of distributions we expect to encounter. This improved understanding can be coupled with further data exploration.

Appendix

The empirical analysis of ranking based datasets depends on the availability of large amounts of data depicting different types of real scenarios. For our experimental setup we used two different datasets from “Rating and Combinatorial Preference Data” of <http://www.preflib.org/data/>. One of these datasets contains 675,069 ratings on scale 1-5 of 1,842 hotels from Trip Advisor website. Other consists of 398 approval ballots and subjective ratings on a 20-point scale collected over 15 potential candidates for the 2002 French Presidential election. The rating was provided by students at Institut d’Etudes Politiques de Paris. The main quantities we are interested in computing are the number of times that each value of k produces the lowest error over the items, and the average value of the errors over all items for each k value. We also provide experimental results from the Jester Online Recommender System on joke ratings.

Trip advisor hotel rating

In the first set of experiments, the dataset contains different types of rating based on the price, quality of rooms, proximity of location, cleanliness, availability of service from front desk etc and finally overall rating provided by the users scraped from Trip advisor. We compared performance between using $k = 2, 3, 4, 5$ to see for how many of the trials each value of k produced the minimal error using the floor approach (Table 11). Surprisingly, we see that the number of victories sometimes decreases with the increase in value of k , while the average error decreased monotonically (recall that we would have zero error if we set k to the actual maximum rating point). The number of victories increases for some cases in $k=2$ vs. 3 compared to 2 vs. 4 (Table 13).

We next explored rounding to generate the rating. For different value of $k \in \{2, 3, 4\}$, all the ratings provided by users was compressed with the computed k midpoints and the average score was calculated. Table 14 shows the average error induced by the compression which performs better than the floor approach for this dataset. An interesting observation found for rounding is that using $k = n = 5$ was outperformed by using $k = 4$ for some of the ratings. Also # of victories was higher in some cases as shown in Table 17.

Average error	k = 2	3	4
Overall	1.04	0.31	0.15
Price	1.07	0.27	0.14
Rooms	1.06	0.32	0.16
Location	1.47	0.42	0.16
Cleanliness	1.43	0.40	0.16
Front Desk	1.34	0.33	0.14
Service	1.24	0.32	0.14
Business Service	0.96	0.28	0.18

Table 11: Average flooring error for hotel ratings.

French presidential election

Voters rated the potential candidates on a scale of 20 where 0.0 is the lowest possible rating and -1.0 indicates missing rating value. For the experiment, we have only considered

Minimal error	k = 2	3	4
Overall	235	450	1157
Price	181	518	1143
Rooms	254	406	1182
Location	111	231	1500
Cleanliness	122	302	1418
Front Desk	120	387	1335
Service	140	403	1299
Business Service	316	499	1027

Table 12: Number of times each k produces minimal error.

# of victories	k= 2 vs. 3	2 vs. 4	3 vs. 4
Overall	243, 1599	277, 1565	5, 1837
Price	187, 1655	211, 1631	4, 1838
Rooms	275, 1567	283, 1559	10, 1832
Location	126, 1716	122, 1720	11, 1831
Cleanliness	126, 1716	141, 1701	5, 1837
Front Desk	130, 1712	133, 1709	8, 1834
Service	153, 1689	152, 1690	11, 1831
Business Service	368, 1474	329, 1513	22, 1820

Table 13: Number of times k value produces minimal error.

Average error	k = 2	3	4
Overall	0.50	0.28	0.15
Price	0.48	0.31	0.15
Rooms	0.48	0.30	0.16
Location	0.63	0.41	0.22
Cleanliness	0.6	0.4	0.21
Front Desk	0.55	0.39	0.21
Service	0.52	0.36	0.18
Business Service	0.39	0.36	0.18

Table 14: Average error using rounding approach.

Minimal error	k = 2	3	4
Overall	82	132	1628
Price	92	74	1676
Rooms	152	81	1609
Location	93	52	1697
Cleanliness	79	44	1719
Front Desk	89	50	1703
Service	102	29	1711
Business Service	246	123	1473

Table 15: Number of times k minimizes error with rounding.

# of victories	k= 2 vs. 3	2 vs. 4	3 vs. 4
Overall	161, 1681	113, 1729	486, 1356
Price	270, 1572	101, 1741	385, 1457
Rooms	344, 1498	173, 1669	575, 1267
Location	275, 1567	109, 1733	344, 1498
Cleanliness	210, 1632	90, 1752	289, 1553
Front Desk	380, 1462	95, 1747	332, 1510
Service	358, 1484	109, 1733	399, 1443
Business Service	870, 972	278, 1564	853, 989

Table 16: Number of times k minimizes error with rounding.

ratings of the candidates. The # of victories and minimal error while using floor were consistent for all comparison

Overall	Average error # of victories	0.15, 0.21 1007, 835
Price	Average error # of victories	0.15, 0.17 955, 887
Rooms	Average error # of victories	0.15, 0.23 1076, 766
Location	Average error # of victories	0.22, 0.22 694, 1148
Cleanliness	Average error # of victories	0.21, 0.19 653, 1189
Front Desk	Average error # of victories	0.21, 0.17 662, 1180
Service	Average error # of victories	0.18, 0.18 827, 1015
Business Service	Average error # of victories	0.18, 0.31 1233, 609

Table 17: # victories and average rounding error, k in $\{4,5\}$.

of k values resulting in higher error for lower value of k for each individual candidate. On the other hand, # of victories and minimal error with rounding compression only produces minimal error for just one of the instance of $k = 2$; the rest were produced by the highest value of k i.e $k = 10$.

Joke recommender system

We also experimented on anonymous ratings data from the Jester Online Joke Recommender System (Goldberg et al. 2001). Data was collected from 73,421 anonymous users between April 1999–May 2003 who have rated 36 or more jokes with ratings of real values ranging from -10.00 to $+10.00$. We included data from 24,983 users in our experiment. Each row of the dataset represents rating from single user. The first column contains the number of jokes rated by a user and the next 100 columns give the ratings for jokes 1–100. Due to space limitation we only experimented on a subset of columns. The average error for columns 5, 7, 8, 13, 15, 16, 17, 18, 19, 20 are reported since they are densely populated. The results are provided in Tables 20 and 21.

Average error	2	3	4	5	8	10
Francois Bayrou	3.18	1.5	0.94	0.66	0.3	0.2
Olivier Besancenot	1.7	0.8	0.5	0.35	0.16	0.1
Christine Boutin	1.15	0.54	0.34	0.24	0.11	0.07
Jacques Cheminade	0.64	0.3	0.19	0.13	0.06	0.04
Jean-Pierre Chevenement	3.69	1.74	1.09	0.77	0.35	0.23
Jacques Chirac	3.48	1.64	1.03	0.72	0.33	0.21
Robert Hue	2.39	1.12	0.7	0.49	0.22	0.14
Lionel Jospin	5.45	2.57	1.61	1.13	0.52	0.33
Arlette Laguiller	2.2	1.04	0.65	0.46	0.21	0.13
Brice Lalonde	1.53	0.72	0.45	0.32	0.14	0.09
Corine Lepage	2.24	1.06	0.67	0.47	0.22	0.14
Jean-Marie Le Pen	0.4	0.19	0.12	0.08	0.04	0.02
Alain Madelin	1.93	0.91	0.57	0.4	0.18	0.12
Noel Mamere	3.68	1.74	1.09	0.77	0.35	0.23
Bruno Maigret	0.31	0.15	0.09	0.06	0.03	0.02

Table 18: Average flooring error for French election.

Average error	2	3	4	5	8	10
Francois Bayrou	1.65	0.73	0.91	0.75	0.48	0.62
Olivier Besancenot	3.88	2.39	2.14	1.7	1.31	1.25
Christine Boutin	3.87	2.39	1.84	1.5	0.9	0.86
Jacques Cheminade	4.34	2.72	2.07	1.65	1.02	0.88
Jean-Pierre Chevenement	1.47	0.65	1.2	0.82	0.55	0.61
Jacques Chirac	1.64	1.0	1.13	0.88	0.55	0.64
Robert Hue	2.51	1.27	1.14	1.09	0.67	0.77
Lionel Jospin	0.33	0.49	0.87	0.67	0.51	0.63
Arlette Laguiller	2.62	1.34	1.34	1.02	0.6	0.63
Brice Lalonde	3.45	1.9	1.55	1.21	0.66	0.78
Corine Lepage	2.89	1.59	1.56	1.16	0.79	0.87
Jean-Marie Le Pen	4.92	3.26	2.55	2.06	1.39	1.2
Alain Madelin	3.18	1.8	1.52	1.17	0.72	0.7
Noel Mamere	2.02	1.55	1.77	1.44	1.29	1.41
Bruno Maigret	4.88	3.23	2.46	1.99	1.28	1.1

Table 19: Average error using rounding approach.

Average error	k=2	3	4	5	10
Joke 5	0.57	0.53	0.52	0.51	0.5
Joke 7	1.32	0.88	0.74	0.66	0.54
Joke 8	1.51	0.97	0.8	0.71	0.56
Joke 13	2.52	1.45	1.09	0.91	0.61
Joke 15	2.48	1.43	1.08	0.91	0.62
Joke 16	3.72	2.01	1.44	1.16	0.69
Joke 17	1.94	1.18	0.92	0.8	0.58
Joke 18	1.51	0.97	0.79	0.71	0.56
Joke 19	0.8	0.64	0.58	0.56	0.51
Joke 20	1.77	1.1	0.87	0.76	0.57

Table 20: For each value of k in $\{2,3,4,5,10\}$, Average error using floor approach, Jester Online Joke Ratings.

Average error	k=2	3	4	5	10
Joke 5	0.48	0.47	0.48	0.47	0.48
Joke 7	1.2	1.2	1.2	1.2	1.2
Joke 8	1.44	1.43	1.42	1.43	1.42
Joke 13	2.43	2.43	2.43	2.42	2.42
Joke 15	2.34	2.34	2.33	2.33	2.33
Joke 16	3.59	3.58	3.57	3.57	3.57
Joke 17	1.84	1.82	1.82	1.81	1.81
Joke 18	1.45	1.44	1.44	1.44	1.44
Joke 19	0.72	0.72	0.71	0.71	0.71
Joke 20	1.65	1.63	1.63	1.63	1.63

Table 21: For each value of k in $\{2,3,4,5,10\}$, Average error using Rounding approach for Jester Online Joke Ratings.

References

- Alwin, D. F. 1997. Feeling thermometers versus 7-point scales: Which are better? *Sociological Methods and Research* 25:318–340.
- Ankenman, J., and Chen, B. 2006. *The Mathematics of Poker*. Pittsburgh, PA, USA: ConJelCo LLC.
- C Preston, C., and Colman, A. 2000. Optimal number of response categories in rating scales: Reliability, validity, discriminating power, and respondent preferences. *Acta psychologica* 104:1–15.
- Cox III, E. P. 1980. The optimal number of response alternatives for a scale: A review. *Journal of Marketing Research* 17:407–442.
- Dubey, P., and Geanakoplos, J. 2010. Grading exams: 100, 99, 98, ... or A, B, C? *Games and Economic Behavior* 69:72–94.
- Friedman, H. H.; Wilamowsky, Y.; and Friedman, L. W. 1981. A comparison of balanced and unbalanced rating scales. *The Mid-Atlantic Journal of Business* 19(2):1–7.
- Fry, H. 2015. *The Mathematics of Love*. New York, NY: TED Books.
- Garland, R. 1991. The mid-point on a rating scale: is it desirable? *Marketing Bulletin* 2:66–70. Research Note 3.
- Goldberg, K.; Roeder, T.; Gupta, D.; and Perkins, C. 2001. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval* 4(2):133–151. <http://eigentaste.berkeley.edu>.
- Lewis, J. R., and Erdiñç, O. 2017. User experience rating scales with 7, 11, or 101 points: Does it matter? *J. Usability Studies* 12(2):73–91.
- Likert, R. A. 1932. A technique for measurement of attitudes. *Archives of Psychology* 22:1–.
- Lozano, L.; Garca-Cueto, E.; and Muiz, J. 2008. Effect of the number of response categories on the reliability and validity of rating scales. *Methodology* 4:73–79.
- Matell, M. S., and Jacoby, J. 1971. Is there an optimal number of alternatives for likert scale items? Study 1: reliability and validity. *Educational and psychological measurements* 31:657–674.
- Matell, M. S., and Jacoby, J. 1972. Is there an optimal number of alternatives for likert scale items? Effects of testing time and scale properties. *Journal of Applied Psychology* 56(6):506–509.
- McKennell, A. C. 1974. Surveying attitude structures: A discussion of principles and procedures. *Quality and Quantity* 7(2):203–294.
- R. Lehmann, D., and Hulbert, J. 1972. Are three-point scales always good enough? *Journal of Marketing Research* 9:444.
- S. Iyengar, S., and Lepper, M. 2001. When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology* 79:995–1006.
- Sudman, S.; Bradburn, N. M.; and Schwartz, N. 1996. *Thinking About Answers: The Application of Cognitive Processes to Survey Methodology*. San Francisco, California: Jossey-Bass Publishers.
- Wildt, A. R., and Mazis, M. B. 1978. Determinants of scale response: label versus position. *Journal of Marketing Research* 15:261–267.