

CrowDEA: Multi-View Idea Prioritization with Crowds

Yukino Baba
 University of Tsukuba
 baba@cs.tsukuba.ac.jp

Jiyi Li
 University of Yamanashi
 jyli@yamanashi.ac.jp

Hisashi Kashima
 Kyoto University
 kashima@i.kyoto-u.ac.jp

Abstract

Given a set of ideas collected from crowds with regard to an open-ended question, how can we organize and prioritize them in order to determine the preferred ones based on preference comparisons by crowd evaluators? As there are diverse latent criteria for the value of an idea, multiple ideas can be considered as “the best”. In addition, evaluators can have different preference criteria, and their comparison results often disagree. In this paper, we propose an analysis method for obtaining a subset of ideas, which we call frontier ideas, that are the best in terms of at least one latent evaluation criterion. We propose an approach, called CrowDEA, which estimates the embeddings of the ideas in the multiple-criteria preference space, the best viewpoint for each idea, and preference criterion for each evaluator, to obtain a set of frontier ideas. Experimental results using real datasets containing numerous ideas or designs demonstrate that the proposed approach can effectively prioritize ideas from multiple viewpoints, thereby detecting frontier ideas. The embeddings of ideas learned by the proposed approach provide a visualization that facilitates observation of the frontier ideas. In addition, the proposed approach prioritizes ideas from a wider variety of viewpoints, whereas the baselines tend to use to the same viewpoints; it can also handle various viewpoints and prioritize ideas in situations where only a limited number of evaluators or labels are available.

1 Introduction

Despite the recent advances in artificial intelligence, there are still several challenges that humans can handle better than machines, especially abstract, open-ended, and context-dependent problems. Brainstorming new ideas is a typical example; for instance, to answer open-ended questions, such as “What is the best logo for the next summer Olympic games?”, “How can we reduce the number of latecomers at team meetings”, and “What are the most reasonable solutions for preventing global warming?”, humans are expected to present more creative and reasonable solutions than machines. Existing studies demonstrate that crowdsourcing is an effective approach to collecting several creative ideas from a wide range of people (Yu and Nickerson 2011;

Koyama, Sakamoto, and Igarashi 2014; Siangliulue et al. 2015; Prpić et al. 2015).

Let us consider the example of designing a suitable logo for the next Olympic games. For example, let us assume that we ask crowd workers to provide a set of candidate designs. After collecting several design ideas, we should organize and prioritize them to select the best. However, the criteria for the best design are usually multi-faceted; for example, there may be two different criteria for design, e.g., traditional aesthetics and contemporary aesthetics. Therefore, there rarely exists a single overwhelming winner over the other candidates in terms of all criteria. Moreover, it is often difficult to define the criteria in advance.

Thus, we must turn to the crowd for assistance, with the expectation that crowd evaluators may be able to identify the unknown diverse criteria. We must ask them to evaluate the ideas, often in the form of pairwise preference comparisons. The criteria for these comparisons can also be diverse depending on evaluators’ personal viewpoints.

In this study, we consider the problem of aggregating the pairwise idea preference comparisons by crowds containing different viewpoints so that a set of best ideas from certain viewpoints may be obtained. These ideas are called *frontier ideas*. The proposed method, which is called CROWDEA, generates a priority map that is a low-dimensional latent space, where ideas are embedded such that the frontier ideas are furthest from the origin and the ideas projected onto the viewpoint of each evaluator are consistent with their pairwise comparisons.

Existing studies (Bradley and Terry 1952; Causeur and Husson 2005; Chen et al. 2013) estimate a unique rank list from the pairwise preference comparisons; they usually assume that there exists a unique rank list as the ground truth. In addition, as there are no explicit evaluation criteria readily available, existing methods, such as skyline query (Borzsony, Kossmann, and Stocker 2001; Hose and Vlachou 2012; Lofi, El Maarry, and Balke 2013), cannot be used. The priority map of CROWDEA assists in making the final decision or further analysis (such as next-round idea sourcing) by providing an organized view from various perspectives.

We provide an illustrative example in Fig. 1; there are




Figure 1: Illustrative example of the proposed multi-view analysis CROWDEA. (a) Target objects have different sizes and colors. (b) Pairwise comparison is performed by crowd evaluators with individual preferences. (c) CROWDEA yields a *priority map*, which is a multiple-criteria preference space, where the objects are embedded so that promising candidates are found as the *frontier objects*. The largest object and the darkest-colored object as well as the fairly large-and-dark object are on the frontier (shown by the dotted line).

nine objects with different sizes and colors (Fig. 1(a)), and we have to prioritize them in terms of various latent criteria, such as size and color. We ask crowd evaluators to make pairwise preference comparisons based on their own personal criteria (Fig. 1(b)). For example, some evaluators prefer darker objects regardless of the object size, whereas others prefer larger objects. CROWDEA outputs the priority map (Fig. 1(c)), where the frontier objects are placed on the convex hull (shown by the dotted line) of all the embedded objects. The x -axis is interpreted as the object size and the y -axis as the color darkness. The rightmost and topmost objects are the best according to the size and darkness criteria, respectively. In addition, the top-right object is the best in terms of an intermediate criterion. The object is both fairly large and dark-colored, making it also a promising candidate.

We verify the proposed approach using real datasets that contain numerous ideas or designs. The quantitative results and qualitative analysis demonstrate that CROWDEA outperforms the baselines. The contributions of this study are as follows:

- We define a problem that involves organizing and prioritizing a set of ideas from multiple preference viewpoints to support decision-making.
- We propose an approach that prioritizes ideas from multiple viewpoints based on pairwise preference comparisons by crowd evaluators. The proposed approach can effectively determine the frontier ideas in a set of ideas.
- The embeddings of ideas learned by the proposed approach provide a visualization that facilitates observation of the frontier ideas; in addition, the proposed approach prioritizes ideas from a wider variety of viewpoints, whereas the baselines tend to use the same viewpoints. The proposed approach can also handle various viewpoints and prioritize ideas in situations where only a limited number of evaluators or labels are available.

2 Related Work

2.1 Idea Crowdsourcing

Existing studies demonstrate that crowdsourcing is an effective method for collecting several creative ideas from a wide range of people (Yu and Nickerson 2011; Siangliulue et al. 2015; Prpić et al. 2015). To understand a set of ideas, it is important to organize and visualize them. Several studies considered with crowdsourcing for organizing ideas. Siangliulue et al. proposed an idea map to visualize a set of ideas using triple-wise similarity queries (Siangliulue et al. 2015). Ahmed and Fuge proposed to find high quality ideas by using community feedback, idea uniqueness, and text features (Ahmed and Fuge 2017). Li et al. proposed an approach that simultaneously ranks and clusters ideas (Li, Baba, and Kashima 2018). In contrast to these approaches, we allow multiple criteria so that promising candidates can be obtained from various viewpoints (i.e., frontier ideas). Similar to our work, Lykourantzou et al. proposed a strategy for ranking ideas according to quality and diversity (Lykourantzou et al. 2018). In their work, the diversity was measured by using the results of manual clustering although our method does not require such manual effort.

2.2 Decision Support Methods

Mathematical methods for supporting decision making have been traditionally studied in operations research. For example, data envelopment analysis (DEA) is a nonparametric method for estimating production frontiers (Seiford and Thrall 1990; Cooper, Seiford, and Zhu 2004), from which the proposed notion of frontier ideas was inspired. The skyline query method, which retains only the objects that are not worse than any others in terms of at least one evaluation criterion, has been extensively studied (Borzsony, Kossmann, and Stocker 2001; Hose and Vlachou 2012; Lofi, El Maarry, and Balke 2013). In contrast with DEA and skyline query, the proposed frontier analysis does not require explicit evaluation criteria, and latent evaluation criteria are learned from the data.

2.3 Pairwise Preference Aggregation

Methods for aggregating pairwise comparison results have long been discussed. The Bradley–Terry (BT) model (Bradley and Terry 1952) is a well-known model for pairwise comparisons. It estimates a single competency score for each object so that the scores are consistent with the pairwise comparison labels. To model more complex object relationships, multi-dimensional generalizations of the BT model have been proposed, such as, the multi-dimensional BT model (Causeur and Husson 2005) and intransitivity model (Chen and Joachims 2016a; 2016b; Duan et al. 2017). The BT model has also been extended to allow variability in the evaluators (Chen et al. 2013). Our work can be considered as the intersection of the above two extensions; we consider multi-dimensional criteria for both evaluators and evaluated objects.

2.4 Multi-View Representation

In some studies on learning multi-view representations, the term ‘multi-view’ has multiple meanings. In several cases, it implies that data instances are described by different types of explicit features (Li, Yang, and Zhang 2016; Wang et al. 2015), for example, images and texts (Li, Yang, and Zhang 2016), texts in two different languages (Chandar et al. 2014), and audio and video media (Huang and Kingsbury 2013). Amid and Ukkonen targeted multiple implicit attributes, where object similarity from triple-wise questions is preserved (Amid and Ukkonen 2015). Their goal is to obtain a space reflecting object similarity, whereas we obtain a space reflecting idea priority.

2.5 Personalized Ranking

Personalized ranking in recommendation systems in which the relative preference of each user is estimated has been extensively studied. For example, Rendle et al. proposed Bayesian personalized ranking, which trains a matrix factorization model to optimize a ranking loss function (Rendle et al. 2009). This topic has been studied in various scenarios, such as group preference (Pan and Chen 2013), visual recommendation (He and McAuley 2016), and event recommendation (Qiao et al. 2014). Their focus is on predicting personalized sets of items for different users, whereas we are interested in obtaining the most advantageous evaluation criterion for each item so that all promising items (i.e., ideas) for decision making may be determined. This results in a different formulation.

2.6 Search Result Diversification

When using web search, users expect not only the most relevant search results to a given query but also diverse ones. Some studies provide both diverse and representative results in terms of content and semantic information (Kennedy and Naaman 2008; Wang et al. 2010), and there are studies on the users’ potential intents (such as navigational or informational) of their queries based on their search behaviors (Cheng, Gao, and Liu 2010; Santos, Macdonald, and Ounis 2011). An important difference between the above-mentioned studies and this one lies in the problem setting:

explicit features such as content or context are not available, and we prioritize ideas based solely on pairwise preferences rather than features. Another difference is that many of these studied have predefined viewpoints, such as the types of user intents, while ours finds the viewpoints from preference comparisons.

3 Multi-View Idea Prioritization with Crowds

3.1 Models and Problem Setting

We address the problem of prioritizing a collection of n ideas in terms of different latent evaluation criteria. Let $[n] = \{1, 2, \dots, n\}$. Then, we consider the embedding \mathbf{x}_i for each idea $i \in [n]$ in a d -dimensional space, which we call the *priority map*. Each axis of the priority map corresponds to a latent preference criterion, and a large value on an axis implies high preference in terms of the corresponding criterion.

Decisions are usually made not only according to a single criterion but also by balancing different criteria. For every idea, there should be a viewpoint that best emphasizes its merits, and it is beneficial to determine the set of all ideas that are “the best” from certain viewpoints. We define the best viewpoint for an idea i as a d -dimensional unit vector \mathbf{v}_i , where the projection of \mathbf{x}_i onto \mathbf{v}_i (i.e., $\mathbf{v}_i^\top \mathbf{x}_i$) is considered to be its preference score from that viewpoint. If idea i is the most preferred among all the ideas, i.e., $\mathbf{v}_i^\top \mathbf{x}_i > \mathbf{v}_i^\top \mathbf{x}_j$, for all $j \neq i$, the idea is promising and should be further investigated. The goal is to determine these ideas, which we call *frontier ideas*; they are located on the convex hull (indicated by the dotted line in Fig. 1(c)) of all ideas in the embedding space. It should be noted that not all ideas can be the best, even from their best viewpoints.

To create the priority map, we collect preference data from m crowd evaluators in the form of pairwise comparisons. Let $\mathcal{C}_k = \{(i, j) \mid i, j \in [n], i \succ_k j\}$ be the set of pairwise comparison results by evaluator $k \in [m]$, where $i \succ_k j$ indicates that evaluator k prefers idea i over idea j . As in the case of the best viewpoints for ideas, every crowd evaluator has its individual viewpoint. We define the viewpoint of crowd evaluator k as a d -dimensional unit vector, \mathbf{w}_k . The projections of $\{\mathbf{x}_i\}_{i=1}^n$ onto \mathbf{w}_k , i.e., $\{\mathbf{w}_k^\top \mathbf{x}_i\}_{i=1}^n$, are regarded as the preference scores by the evaluator, and they are expected to be consistent with the pairwise comparison results, \mathcal{C}_k .

In summary, the inputs and outputs of the problem are as follows:

Inputs: n ideas, m crowd evaluators, and $\{\mathcal{C}_k\}_{k=1}^m$, where $\mathcal{C}_k = \{(i, j) \mid i, j \in [n], i \succ_k j\}$ is the set of pairwise comparison results by evaluator $k \in [m]$.

Outputs: $\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{v}_i\}_{i=1}^n, \{\mathbf{w}_k\}_{k=1}^m$, where \mathbf{x}_i is the d -dimensional embedding of idea $i \in [n]$, \mathbf{v}_i is the best viewpoint for idea $i \in [n]$, and \mathbf{w}_k is the viewpoint of crowd evaluator $k \in [m]$.

3.2 Estimation

We formulate the multi-view analysis as an optimization problem. Based on the discussions in the previous section,

we have two optimization sub-goals: (i) determine as many frontier ideas as possible, and (ii) achieve consistency with the pairwise preference comparison results.

For the first sub-goal, we impose the best viewpoint for each idea, from which the idea is most valuable among all ideas. That is, we require that the resultant idea embeddings $\{\mathbf{x}_i\}_{i=1}^n$ and corresponding best viewpoints $\{\mathbf{v}_i\}_{i=1}^n$ satisfy the constraints

$$\mathbf{v}_i^\top \mathbf{x}_i > \mathbf{v}_i^\top \mathbf{x}_j, \forall i \in [n], \forall j \neq i \in [n]. \quad (1)$$

As it is not possible to satisfy all of the constraints, we quantify the number of constraint violations using a loss function. Specifically, we use the hinge loss as the loss function:

$$\mathcal{L}_F(\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{v}_i\}_{i=1}^n) = \frac{1}{n(n-1)} \sum_{i \in [n]} \sum_{j \in [n] \setminus i} \max\{0, 1 - \mathbf{v}_i^\top (\mathbf{x}_i - \mathbf{x}_j)\}. \quad (2)$$

For the second sub-goal, the aim is to make the viewpoint of each evaluator consistent with the pairwise comparison results by that evaluator. We assume that each crowd evaluator has their own viewpoint, and we define \mathbf{w}_k as the preference criterion vector for the preference labels of evaluator k . From the viewpoint of evaluator k , the preference score of each idea i is given as $\mathbf{w}_k^\top \mathbf{x}_i$; therefore, the set \mathcal{C}_k of all pairwise comparison results by evaluator k should be consistent with the preference scores, i.e.,

$$\mathbf{w}_k^\top \mathbf{x}_i > \mathbf{w}_k^\top \mathbf{x}_j, \forall k \in [m], \forall (i, j) \in \mathcal{C}_k. \quad (3)$$

As before, it is not always possible to meet all of the constraints, and again we use the hinge loss function:

$$\mathcal{L}_C(\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{w}_k\}_{k=1}^m) = \frac{1}{c} \sum_{k \in [m]} \sum_{i, j \in \mathcal{C}_k} \max\{0, 1 - \mathbf{w}_k^\top (\mathbf{x}_i - \mathbf{x}_j)\}, \quad (4)$$

where $c = \sum_k |\mathcal{C}_k|$ is the number of observed preference labels.

In addition, we impose the constraint that all embeddings and preference criterion vectors should be non-negative for a more intuitive visualization (as shown in Fig. 1(c)). Furthermore, we add the constraints that all the preference criterion vectors, \mathbf{w}_k and \mathbf{v}_i , have unit length, i.e., $\|\mathbf{w}_k\|_2 = 1$ and $\|\mathbf{v}_i\|_2 = 1$. One advantage of this constraint is that it scales the embeddings for all objects. This unit length constraint can also avoid the preference criterion vector being zero. For example, for an object o_i that is not on the frontier and ranked low even in its best viewpoint, if \mathbf{v}_i is not equal to zero, $\mathbf{v}_i^\top (\mathbf{x}_i - \mathbf{x}_j)$ for many o_j are lower than zero, which may result in $\mathbf{v}_i = \mathbf{0}$ minimizing $\mathcal{L}_F(\mathbf{x}_i, \mathbf{v}_i)$.

By combining the loss functions for the two sub-goals and the constraints, the optimization problem can be fully formulated as follows:

$$\begin{aligned} & \underset{\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{v}_i\}_{i=1}^n, \{\mathbf{w}_k\}_{k=1}^m}{\text{minimize}} && \mathcal{L}_C(\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{w}_k\}_{k=1}^m) \\ & && + \alpha \mathcal{L}_F(\{\mathbf{x}_i\}_{i=1}^n, \{\mathbf{v}_i\}_{i=1}^n) \\ \text{subject to} &&& \mathbf{x}_i, \mathbf{v}_i, \mathbf{w}_k \in \mathbb{R}_+^d, \forall i \in [n], k \in [m]; \\ &&& \|\mathbf{w}_k\|_2 = 1, \forall k \in [m]; \\ &&& \|\mathbf{v}_i\|_2 = 1, \forall i \in [n], \end{aligned}$$

where $\alpha > 0$ is a constant that controls the trade-off between \mathcal{L}_C and \mathcal{L}_F .

The constrained optimization is performed in a straightforward fashion; after the optimization algorithm updates the parameters at each step, all negative entries are set to zero to satisfy the non-negativity constraints; each \mathbf{w}_k and \mathbf{v}_i is then normalized to satisfy the unit length constraint. Finally, idea i is considered a frontier idea if there exists \mathbf{v} that satisfies $\|\mathbf{v}\|_2 = 1$, $\mathbf{v} \geq 0$, and $\mathbf{v}^\top \mathbf{x}_i > \mathbf{v}^\top \mathbf{x}_j$ for all $j \neq i \in [n]$.

4 Experiments

4.1 Experimental Design

We empirically evaluate the proposed method using real datasets containing ideas and designs for pairwise comparison. The experiments were designed to answer the following questions:

- Q1. Visualization:** How successful is CROWDEA in organizing ideas?
- Q2. Accuracy:** How accurately does CROWDEA prioritize ideas according to multiple viewpoints?
- Q3. Efficiency:** How does the accuracy change according to the number of evaluators?

4.2 Datasets

We constructed two types of real datasets (Table 1 summarizes the data statistics)¹:

- **Ideas:** We prepared five open-ended day-to-day life questions, such as ‘‘How can we reduce the number of late-comers at team meetings?’’, and we collected solution ideas from crowdsourcing workers using the crowdsourcing platform, Lancers. We obtained approximately 80 ideas for each question. We hired another set of crowd workers for collecting preference labels, and we asked them to compare pairs of ideas for each problem. Approximately 20 workers were assigned for each pair of ideas, and each worker evaluated at least 50 pairs. The order of pairs and that of ideas in each pair were randomized. There were approximately 160–260 evaluators and 64K preference labels in total for each dataset.
- **Designs:** We held a character design contest for an artificial intelligence (AI) research laboratory and collected 66 designs. We also prepared 38 logos for the summer and winter Olympic games from 1948 to 2020, and we collected preference labels for these two design tasks in the same manner as for the datasets containing ideas. There were 183 evaluators and 43K preference labels for the ‘‘Character’’ dataset and 64 evaluators and 14K labels for the ‘‘Olympic’’ dataset.

4.3 Baselines

We compare CROWDEA with the following four baselines (They are summarized in Table 2):

¹Datasets, codes, and Jupyter notebook for reproducing tables and figures are available at: <https://github.com/yukinobaba/crowdea>.

Table 1: Summary dataset statistics

(a) Ideas				
Dataset	Problem	#ideas	#evaluators	#labels
Bike	“How can we discourage indiscriminate bicycle parking on campus?”	81	217	64,800
Cheat	“How can we effectively prevent students from cheating in exams?”	80	257	63,200
Meeting	“How can we reduce the number of latecomers for team meetings?”	80	177	63,200
Night	“How can we stay safe when walking alone at night?”	80	171	63,200
Visitor	“How can we support foreign tourists who encounter a language barrier?”	81	158	64,800

(b) Designs				
Dataset	Problem	#ideas	#evaluators	#labels
Olympics	“Design a logo for the Olympic Games.”	38	64	14,100
Character	“Design a character for an AI research laboratory.”	66	183	42,928

Table 2: Comparison of CrowDEA and baselines

	Multi-evaluators	Multi-dimensional	Multi-view
BT	-	-	-
CROWDBT	✓	-	-
BLADE-CHEST	-	✓	-
BPR	✓	✓	-
CROWDEA	✓	✓	✓

- **BT** (Bradley and Terry 1952) is the Bradley–Terry (BT) model, a standard approach for aggregating pairwise preferences. This model represents a preference score for each item by a scalar value and does not assume a different viewpoint for each evaluator.
- **CROWDBT** (Chen et al. 2013) is an extension of BT that incorporates the diversity of evaluator reliability into the model.
- **BLADE-CHEST** (Chen and Joachims 2016a) is a multi-dimensional extension of BT and it models intransitivity in pairwise preference.
- **BPR** (Rendle et al. 2009) is a method for recommendation, which models both the item embedding and user preference by using d -dimensional vectors.

The regularization parameter of the baseline methods was chosen from $\{0.001, 0.01, 0.1\}$, and the best case for a target metric is presented in the results. Although there exist several related studies, most of them are not applicable to the present problem setting; only the results of pairwise comparison are given, whereas the features of each idea are unavailable.

4.4 Setup

α was set to 0.1 in all experiments to achieve a good balance between \mathcal{L}_C and \mathcal{L}_F . If α is large, \mathcal{L}_F pushes all ideas to the frontier, which does not promote detecting the best ideas, whereas a small α lets the frontier ideas form a small and meaningful subset. As the proposed method aims to generate priority maps, we set $d = 2$ or $d = 3$.

4.5 Q1: Visualization

We conducted a case study with design datasets to investigate how well CROWDEA visually organizes the ideas from multiple viewpoints. We applied CROWDEA (with $d = 2$) to all the preference labels in the dataset, and the estimated two-dimensional embeddings were used for generating the priority map shown in Fig. 2a. It can be observed that CROWDEA organizes the ideas along with the frontier curve; CROWDEA can locate each idea receiving a higher preference score (from its best viewpoint), and the priority map thus shows the frontier curve. This provides a well-organized visualization, which facilitates the evaluation of ideas from multiple viewpoints. As mentioned in the introduction, the priority map created by CROWDEA allows us to recognize a variety of viewpoints, such as contemporary aesthetics (x -axis) and traditional aesthetics (y -axis). Recent Olympic logos are placed in the bottom-right region, whereas older logos from the '60s to '80s are placed in the upper-left region, which possibly correlates with the ages of those who provide the preference labels. It should be noted that the above interpretations of the axes are not given in advance. In the priority map, Nagano (1998) Olympics² and Calgary (1988) Olympics, which are highlighted in red, are the two winners on each of the two axes. The priority maps can also capture combinations of these two perspectives, and the winners on them are highlighted in blue in Fig. 2a. Fig. 2b shows the visualization produced by BPR, which achieves the highest accuracy, as presented in Sec. 4.6. In contrast to CROWDEA, BPR assigns much higher priorities to modern logos than traditional ones, and it thus does not produce a frontier curve.

4.6 Q2: Accuracy

We demonstrate how accurately CROWDEA determines the best ideas in various viewpoints.

Setup: We prepared the ground truth of idea priorities from various viewpoints to investigate accuracy. We first

²Nagano (1998) is actually regarded as the best use of athletic imagery by some professional critics, <https://en.99designs.jp/blog/famous-design/olympic-logos/>




Figure 2: Priority maps for the ‘‘Olympic’’ dataset generated by CROWDEA and BPR. CROWDEA produces well-organized visualization and detects good ideas in diverse viewpoints. The top-right corner of each image corresponds to its embedding in the space. The frontier objects detected by CROWDEA are highlighted in red or blue. Both CROWDEA and BPR locate ideas with higher priorities further from the origin. In contrast to BPR, CROWDEA assigns high priorities to the ideas from multiple viewpoints and organizes the ideas along with the frontier curve.

Table 3: Examples of frontier ideas for ‘‘Cheat’’ problem found by CROWDEA. CROWDEA finds worthy ideas in various viewpoints.

Impose severe penalties for cheating, such as cancellation of modules for an entire year.
Prepare two types of examination sheets with differently ordered items, and distribute one to every student such that neighboring students have different exam sheets.
Have proctors watch students from the back of an examination room.
Instead of multiple-choice questions or short answer questions, use essay questions to make it difficult to copy the answers of other students.

collected 100 viewpoints for each dataset from crowdsourcing workers who were shown a pair of ideas and asked to describe a viewpoint that distinguishes the two ideas. For instance, we obtained ‘‘This idea can be easily implemented’’ as a viewpoint for the ‘‘Cheat’’ problem. We then asked workers to grade each idea in terms of each viewpoint on a five-point scale. Ten workers were assigned to each idea-viewpoint pair, and the average grade was used as the ground truth priority p_{ij}^* of idea i from viewpoint j . We removed overlapped or less popular viewpoints by applying k -means clustering to the obtained priorities; that is, we considered $\mathbf{p}_j^* = (p_{1j}^*, \dots, p_{n_j}^*)$ to be the feature vector of viewpoint j and used it for clustering. The number of clusters was set to 50^3 . The clusters with only one sample were then omitted, and the number of remaining clusters was 15–30. We chose the viewpoint closest to the center of each of the remaining clusters, referred to as a representative viewpoint. We thus had 15–30 representative viewpoints for each dataset.

³The representative viewpoints were almost the same when the number of clusters was chosen from $\{30, 40, 50, 60, 70\}$.

We note that neither the proposed method nor the baseline methods can access the ground truth; it is used only for evaluation.

We applied CROWDEA, BPR, and BLADE-CHEST to the preference labels in each dataset and obtained the embeddings $\{\mathbf{x}_i\}_{i=1}^n$. We intended to use the embeddings to rank ideas according to each representative viewpoint in the ground truth to evaluate the ranking accuracy. Given a viewpoint vector \mathbf{v} , the projection of \mathbf{x}_i onto \mathbf{v} (i.e., $\mathbf{v}^\top \mathbf{x}_i$) is considered as the priority score for this viewpoint. We optimize a viewpoint vector \mathbf{v}_j^* , which well represents viewpoint j , according to an evaluation measure. This yielded $p_{ij} = \mathbf{v}_j^{*\top} \mathbf{x}_i$, which is the predicted priority score for that viewpoint. We also applied BT and CROWDBT to the preference labels, and regarded the estimated score p_i as p_{ij} for each viewpoint j . Each method generated a ranking of the ideas for viewpoint j according to $\{p_{ij}\}_i$. We compared these with the ranking by the ground truth priorities, $\{p_{ij}^*\}_i$, and evaluated the ranking accuracy.

The ranking accuracy (nDCG@ k) for each viewpoint was calculated as follows: we had the top k ideas according to the predicted priorities, and their true priorities, $\mathbf{y} = (y_1, \dots, y_k)$, where y_i is the true priority of the i -th ranked idea. We additionally had the true top k ideas and their true priorities $\mathbf{t} = (t_1, \dots, t_k)$. We calculated $\text{DCG}(k, \mathbf{y}) = \sum_{i=1}^k y_i / \log_2(i+1)$ and $\text{IDCG}(k, \mathbf{t}) = \sum_{i=1}^k t_i / \log_2(i+1)$ to obtain $\text{nDCG}@k = \text{DCG}(k, \mathbf{y}) / \text{IDCG}(k, \mathbf{t})$.

Results: Table 3 lists examples of the frontier ideas obtained by CROWDEA for the ‘‘Cheat’’ dataset. It can be seen that CROWDEA provides useful ideas that are considered good from various viewpoints. Table 4 shows the average nDCG@5 and nDCG@10 over the representative viewpoints. It can be seen that CROWDEA outperforms the baselines in most cases; CROWDEA can capture the diversity of viewpoints that are not considered by the other simple methods. Moreover, CROWDEA with $d = 3$ achieves higher scores than with $d = 2$ in all datasets, as the higher-dimensional embedding handles various viewpoints.

Table 4: Average of nDCG@5 and nDCG@10 scores among the representative viewpoints. CROWDEA accurately ranks the ideas according to various viewpoints. The cases in which CROWDEA outperforms the baselines are bold-faced. The cases in which CROWDEA is the statistically significant ($p < 0.05$) winner by the Wilcoxon signed rank test are underlined.

(a) $d = 2$										
Dataset	nDCG@5					nDCG@10				
	BT	CROWDBT	BLADE-CHEST	BPR	CROWDEA	BT	CROWDBT	BLADE-CHEST	BPR	CROWDEA
Bike	<u>0.772</u>	<u>0.779</u>	<u>0.757</u>	0.827	0.833	<u>0.798</u>	<u>0.800</u>	<u>0.756</u>	0.847	0.849
Cheat	<u>0.768</u>	<u>0.767</u>	<u>0.813</u>	<u>0.789</u>	0.893	<u>0.795</u>	<u>0.791</u>	<u>0.819</u>	<u>0.800</u>	0.895
Meeting	<u>0.817</u>	<u>0.815</u>	<u>0.829</u>	<u>0.800</u>	0.877	<u>0.824</u>	<u>0.825</u>	<u>0.837</u>	<u>0.818</u>	0.880
Night	<u>0.790</u>	<u>0.790</u>	<u>0.903</u>	<u>0.853</u>	0.917	<u>0.809</u>	<u>0.808</u>	<u>0.901</u>	<u>0.862</u>	0.912
Visitor	<u>0.818</u>	<u>0.825</u>	<u>0.868</u>	<u>0.933</u>	0.938	<u>0.832</u>	<u>0.835</u>	<u>0.874</u>	<u>0.938</u>	0.943
Character	<u>0.902</u>	<u>0.912</u>	<u>0.866</u>	<u>0.929</u>	0.930	<u>0.911</u>	<u>0.921</u>	<u>0.865</u>	<u>0.926</u>	0.935
Olympic	<u>0.926</u>	<u>0.926</u>	<u>0.920</u>	0.940	0.936	<u>0.937</u>	<u>0.937</u>	<u>0.923</u>	0.949	0.947

(b) $d = 3$										
Dataset	nDCG@5					nDCG@10				
	BT	CROWDBT	BLADE-CHEST	BPR	CROWDEA	BT	CROWDBT	BLADE-CHEST	BPR	CROWDEA
Bike	<u>0.772</u>	<u>0.779</u>	<u>0.803</u>	<u>0.819</u>	0.883	<u>0.798</u>	<u>0.800</u>	<u>0.797</u>	<u>0.835</u>	0.893
Cheat	<u>0.768</u>	<u>0.767</u>	<u>0.847</u>	<u>0.795</u>	0.924	<u>0.795</u>	<u>0.791</u>	<u>0.839</u>	<u>0.804</u>	0.927
Meeting	<u>0.817</u>	<u>0.815</u>	<u>0.867</u>	<u>0.888</u>	0.920	<u>0.824</u>	<u>0.825</u>	<u>0.862</u>	<u>0.891</u>	0.923
Night	<u>0.790</u>	<u>0.790</u>	<u>0.907</u>	<u>0.913</u>	0.953	<u>0.809</u>	<u>0.808</u>	<u>0.894</u>	<u>0.910</u>	0.945
Visitor	<u>0.818</u>	<u>0.825</u>	<u>0.916</u>	<u>0.842</u>	0.955	<u>0.832</u>	<u>0.835</u>	<u>0.906</u>	<u>0.846</u>	0.951
Character	<u>0.902</u>	<u>0.912</u>	<u>0.905</u>	<u>0.957</u>	0.960	<u>0.911</u>	<u>0.921</u>	<u>0.891</u>	0.954	0.953
Olympic	<u>0.926</u>	<u>0.926</u>	<u>0.927</u>	<u>0.956</u>	0.966	<u>0.937</u>	<u>0.937</u>	<u>0.926</u>	<u>0.952</u>	0.964

We quantitatively investigate the variety of the ideas prioritized by the proposed method. Fig. 3 shows the top-10 ideas and a heatmap of the ground truth priority p_{ij}^* of each top-10 idea for each viewpoint. The top-10 ideas ranked by CROWDBT (with $\lambda = 0.01$) and BT ($\lambda = 0.01$) are selected by using p_i , and those by CROWDEA (with $d = 2$) are according to $p_i = \sum_{j \in [n] \setminus i} \mathbf{v}_i^\top (\mathbf{x}_i - \mathbf{x}_j)$, which indicates how likely the ideas are to be frontier ideas. It is observed that CROWDEA prioritizes ideas from a wider variety of viewpoints, whereas the baselines tend to use the same viewpoints. Note that BLADE-CHEST and BPR cannot output a single priority score due to the absence of \mathbf{v}_i .

4.7 Q3: Efficiency

Each dataset contains the preference labels from approximately 200 evaluators; however, it is not always feasible to collect these labels from a large group of evaluators. To demonstrate the efficiency of the proposed method, we evaluate the accuracy of CROWDEA in terms of the number of evaluators. Additionally, each dataset contains 200–400 labels per evaluator. We evaluate the accuracy of CROWDEA in cases where the number of available labels is limited.

Setup: We randomly chose $q \in \{20, 50, 100\}$ evaluators or $r \in \{1000, 2000, 5000, 10000, 20000\}$ labels and applied CROWDEA to the preference labels (i.e., a subset of the preference labels in a dataset). For each q or r , we performed 10 trials and selected a different set of evaluators (or labels) for each trial.

Results: Fig. 4a shows the average nDCG@5 of each method according to the number of evaluators used for model inference. The average nDCG@5 scores are shown

for different viewpoints and ten different subsets of evaluators. The performance of CROWDEA declines as the number of evaluators decreases; however, the average nDCG@5 scores are still over 0.8 in all cases, even when the number of evaluators is only 20, and CROWDEA outperforms the baselines in all cases. Fig. 4b shows the average nDCG@5 of each method according to the number of labels. CROWDEA shows better performance than the other methods even when the number of labels is small. It is worth noting that CROWDEA can handle various viewpoints and prioritize ideas in situations where only a limited number of evaluators or labels are available.

5 Conclusions

We addressed the problem of idea prioritization with crowds. The proposed method estimates the best viewpoint for every idea and preference criterion of every crowd evaluator. Experimental results based on real datasets containing ideas demonstrated that the proposed approach effectively prioritizes ideas from multiple viewpoints and obtains frontier ideas. The visualization based on the learned embeddings facilitates observation of the frontier ideas. Possible future work may include extensions to multiple best viewpoints for each idea, as the present formulation allows only a single best viewpoint. The interpretation of the obtained results is also an important issue; although this is left to users in the present study, systematic interpretation by crowds is an interesting future research direction.

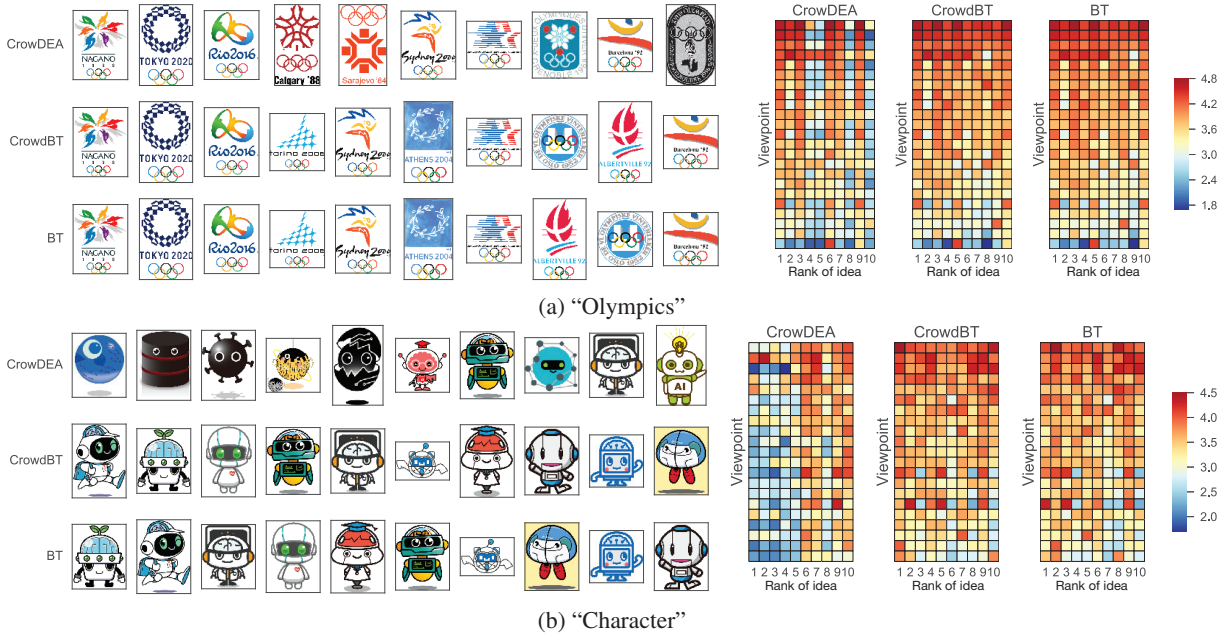


Figure 3: (Left) The top-10 ideas prioritized by each method. The ideas are ordered from left to right according to their estimated preference scores. (Right) The ground truth priority of each of the top-10 ideas in each representative viewpoint. The ideas selected by CROWDEA are prioritized in different viewpoints, while those chosen by the baselines are prioritized in the same viewpoints.

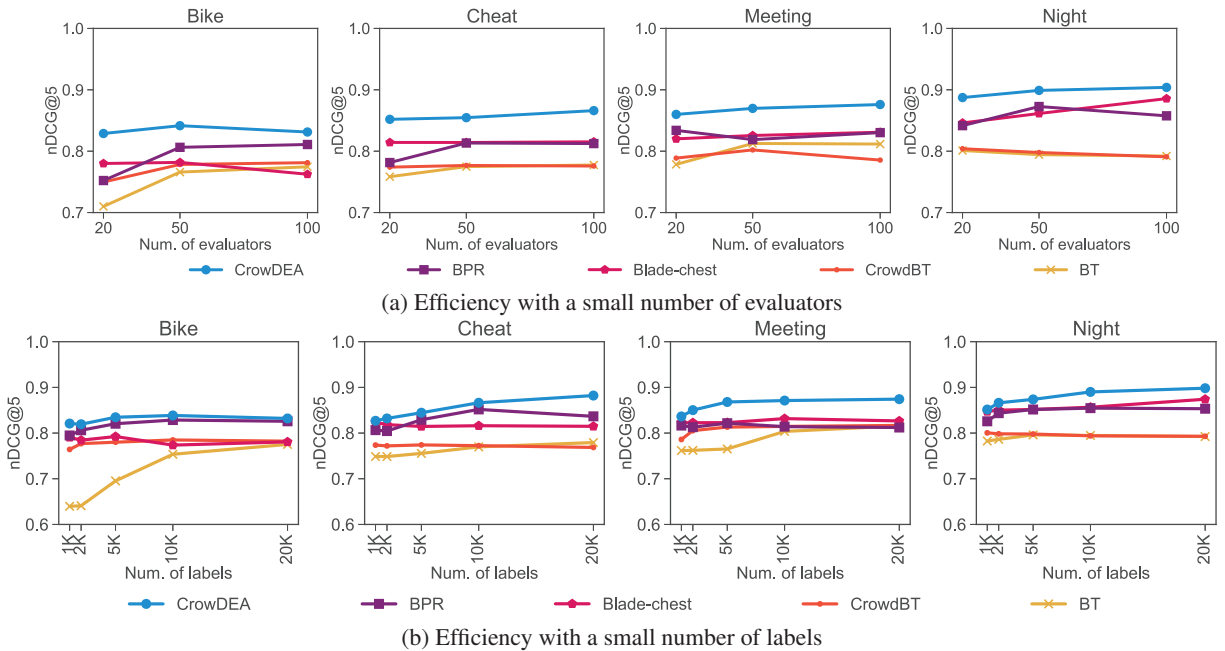


Figure 4: Average of nDCG@5 scores for the representative viewpoints and ten trials. CROWDEA accurately ranks the ideas even when the number of evaluators or the number of labels is small. d is set to 2. Due to space limitation, we only present the results of the first four datasets.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP18K18105 and JST PRESTO Grant Number JP-MJPR19J9, Japan.

References

- Ahmed, F., and Fuge, M. 2017. Capturing winning ideas in online design communities. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*, 1675–1687.
- Amid, E., and Ukkonen, A. 2015. Multiview triplet embedding: learning attributes in multiple maps. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 1472–1480.
- Borzsony, S.; Kossmann, D.; and Stocker, K. 2001. The skyline operator. In *Proceedings of the 17th International Conference on Data Engineering (ICDE)*, 421–430.
- Bradley, R. A., and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika* 39(3/4):324–345.
- Causeur, D., and Husson, F. 2005. A 2-dimensional extension of the bradley–terry model for paired comparisons. *Journal of Statistical Planning and Inference* 135(2):245–259.
- Chandar, S.; Lauly, S.; Larochelle, H.; Khapra, M.; Ravindran, B.; Raykar, V. C.; and Saha, A. 2014. An autoencoder approach to learning bilingual word representations. In *Advances in Neural Information Processing Systems 27*, 1853–1861.
- Chen, S., and Joachims, T. 2016a. Modeling intransitivity in matchup and comparison data. In *Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM)*, 227–236.
- Chen, S., and Joachims, T. 2016b. Predicting matchups and preferences in context. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 775–784.
- Chen, X.; Bennett, P. N.; Collins-Thompson, K.; and Horvitz, E. 2013. Pairwise ranking aggregation in a crowdsourced setting. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining (WSDM)*, 193–202.
- Cheng, Z.; Gao, B.; and Liu, T.-Y. 2010. Actively predicting diverse search intent from user browsing behaviors. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*, 221–230.
- Cooper, W. W.; Seiford, L. M.; and Zhu, J. 2004. Data envelopment analysis. In *Handbook on Data Envelopment Analysis*. Springer. 1–39.
- Duan, J.; Li, J.; Baba, Y.; and Kashima, H. 2017. A generalized model for multidimensional intransitivity. In *Proceedings of the 21st Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 840–852.
- He, R., and McAuley, J. 2016. VBPR: Visual Bayesian personalized ranking from implicit feedback. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 144–150.
- Hose, K., and Vlachou, A. 2012. A survey of skyline processing in highly distributed environments. *The VLDB Journal* 21(3):359–384.
- Huang, J., and Kingsbury, B. 2013. Audio-visual deep learning for noise robust speech recognition. In *Proceedings of 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7596–7599.
- Kennedy, L. S., and Naaman, M. 2008. Generating diverse and representative image search results for landmarks. In *Proceedings of the 17th International Conference on World Wide Web (WWW)*, 297–306.
- Koyama, Y.; Sakamoto, D.; and Igarashi, T. 2014. Crowd-powered parameter analysis for visual design exploration. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST)*, 65–74.
- Li, J.; Baba, Y.; and Kashima, H. 2018. Simultaneous clustering and ranking from pairwise comparisons. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 1554–1560.
- Li, Y.; Yang, M.; and Zhang, Z. 2016. Multi-view representation learning: A survey from shallow methods to deep methods. *arXiv preprint arXiv:1610.01206*.
- Lofi, C.; El Maarry, K.; and Balke, W.-T. 2013. Skyline queries in crowd-enabled databases. In *Proceedings of the 16th International Conference on Extending Database Technology (EDBT)*, 465–476.
- Lykourentzou, I.; Ahmed, F.; Papastathis, C.; Sadien, I.; and Papangelis, K. 2018. When crowds give you lemons: Filtering innovative ideas using a diverse-bag-of-lemons strategy. *Proceedings of the ACM Human Computer Interaction*.
- Pan, W., and Chen, L. 2013. GBPR: Group preference based bayesian personalized ranking for one-class collaborative filtering. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2691–2697.
- Prpić, J.; Shukla, P. P.; Kietzmann, J. H.; and McCarthy, I. P. 2015. How to work a crowd: Developing crowd capital through crowdsourcing. *Business Horizons* 58(1):77–85.
- Qiao, Z.; Zhang, P.; Zhou, C.; Cao, Y.; Guo, L.; and Zhang, Y. 2014. Event recommendation in event-based social networks. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 3130–3131.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI)*, 452–461.
- Santos, R. L.; Macdonald, C.; and Ounis, I. 2011. Intent-aware search result diversification. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 595–604.
- Seiford, L. M., and Thrall, R. M. 1990. Recent developments in dea: the mathematical programming approach to frontier analysis. *Journal of Econometrics* 46(1-2):7–38.

Siangliulue, P.; Arnold, K. C.; Gajos, K. Z.; and Dow, S. P. 2015. Toward collaborative ideation at scale: Leveraging ideas from others to generate more creative and diverse ideas. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*, 937–945.

Wang, M.; Yang, K.; Hua, X.-S.; and Zhang, H.-J. 2010. Towards a relevant and diverse search of social images. *IEEE Transactions on Multimedia* 12(8):829–842.

Wang, W.; Arora, R.; Livescu, K.; and Bilmes, J. 2015. On deep multi-view representation learning. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning (ICML)*, 1083–1092.

Yu, L., and Nickerson, J. V. 2011. Cooks or cobblers?: crowd creativity through combination. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 1393–1402.