

An Improved Phrase-based Approach to Annotating and Summarizing Student Course Responses

Wencan Luo[†] Fei Liu[‡] Diane Litman[†]

[†]University of Pittsburgh, Pittsburgh, PA 15260

[‡]University of Central Florida, Orlando, FL 32816

{wencan, litman}@cs.pitt.edu feiliu@cs.ucf.edu

Abstract

Teaching large classes remains a great challenge, primarily because it is difficult to attend to all the student needs in a timely manner. Automatic text summarization systems can be leveraged to summarize the student feedback, submitted immediately after each lecture, but it is left to be discovered what makes a good summary for student responses. In this work we explore a new methodology that effectively extracts summary phrases from the student responses. Each phrase is tagged with the number of students who raise the issue. The phrases are evaluated along two dimensions: with respect to text content, they should be informative and well-formed, measured by the ROUGE metric; additionally, they shall attend to the most pressing student needs, measured by a newly proposed metric. This work is enabled by a phrase-based annotation and highlighting scheme, which is new to the summarization task. The phrase-based framework allows us to summarize the student responses into a set of bullet points and present to the instructor promptly.

1 Introduction

Effective teachers use student feedback to adjust their teaching strategies. Nowadays, in large classes, there is far too much feedback for a single teacher to manage and attend to. If different perspectives in the student feedback could be summarized and pressing issues identified, it would greatly enhance the teachers' ability to make informed choices. In this work we seek to automatically summarize the student course feedback into a set of bullet points. Each bullet point corresponds to a phrase, tagged with the number of students who raise the issue. Our emphasis is on the textual feedback submitted by students after each lecture in response to two reflective prompts (Boud et al., 2013): 1) “Describe what you found most interesting in today’s class” and 2) “Describe what was confusing or needed more detail.” Education researchers have demonstrated that asking students to respond to reflection prompts can improve both teaching and learning (Van den Boom et al., 2004; Menekse et al., 2011). However, summarizing these responses for large classes (e.g., introductory STEM, MOOCs) remains costly, time-consuming, and an onerous task for humans (Mosteller, 1989).

In our prior work, Luo and Litman (2015) (henceforth **L&L**) introduced the task of automatic summarization of student responses. The challenges of this task include 1) high lexical variety, because students tend to use different word expressions to communicate the same or similar meanings (e.g., “bike elements” vs. “bicycle parts”), and 2) high length variety, as the student responses range from a single word to multiple sentences. To tackle the challenges, L&L proposed a *phrase summarization framework* consisting of three stages: phrase extraction, phrase clustering, and phrase ranking. The approach extracts noun phrases from student responses, groups the phrases using a greedy clustering algorithm, and finally selects representative phrases from the clusters using LexRank (Erkan and Radev, 2004).

There are three limitations in the phrase summarization framework. First, noun phrases do not suffice. Other types of phrases such as “how confidence intervals linked with previous topics” are useful and should be allowed. Second, clustering is based on similarity, but similarity of phrases that do not appear in a background corpus (i.e., the corpus used to learn the similarities) cannot be captured in the previous

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

Reflective Prompt Describe what was confusing or needed more detail.	
<p>Student Responses</p> <p>S1: In the age of distributions example, application of qq plot^g was confusing</p> <p>S2: Last problem about normalization^m</p> <p>S3: central limit theorem^y and A And B events example formulas were different. I did not understand that part well</p> <p>S4: Sampling distribution^r was a little bit abstract</p> <p>S5: Q-q plot^g</p> <p>S6: Central Limit Thm^y</p> <p>S7: CLT^y</p> <p>S8: Normal approximation to binomial^b</p> <p>S9: bernoulli random variables</p> <p>S10: The central limit^y and normal approximations^b</p> <p>...</p>	<p>Human Summary 1</p> <ul style="list-style-type: none"> - central limit theorem^y [12] - q-q plot^g [9] - sampling distribution^r [6] - normal approximation^b [5] - normalization (last example)^m [3] <hr/> <p>Human Summary 2</p> <ul style="list-style-type: none"> - central limit theorem [13] - q-q plots [9] - general more explanations/details, better handwriting, move slower [9] - sampling distributions [6] - nothing [6]

Table 1: Example prompt, student responses, and two human summaries. ‘S1’–‘S10’ are student IDs. The summary phrases are each tagged with the number of students who raise the issue (i.e., student supporters). The summary and phrase highlights are manually created by annotators. Phrases that bear the same color belong to the same issue. Each annotator is free to choose his/her color palette. We have only demonstrated the highlights of **Human Summary 1** to avoid overlaying of two sets of colors on student responses. The superscripts of the phrase highlights are imposed by the authors of this paper to differentiate colors when printed in grayscale (y: yellow, g: green, r: red, b: blue, and m: magenta).

setting. Lastly, a greedy clustering algorithm K-medoids (Kaufman and Rousseeuw, 1987) was previously used to group candidate phrases. It ignores global information and may suffer from a “collapsing” effect, which leads to the generation of a large cluster with unrelated items (Basu et al., 2013).

The goal of this work is to explore a phrase-based highlighting scheme, which is new to the summarization task. We aim to improve the phrase summarization framework by exploiting new capabilities that are enabled by the highlighting scheme. In the new scheme, human annotators are instructed to 1) create summary phrases from the student responses, 2) associate a number with each summary phrase which indicates the number of students who raise the issue (henceforth **student supporters**), and 3) highlight the corresponding phrases in both the human summary and student responses. Table 1 illustrates the highlighting scheme and more details are presented in §3. The new highlighting scheme makes it possible to develop a supervised candidate phrase extraction model (§4.1) and estimate pairwise phrase similarity with supervision (§4.2). To solve the third limitation, we explore a community detection algorithm OSLOM (Lancichinetti et al., 2011) that optimizes the statistical significance of clusters with respect to a global null model (§4.3). Experimental results show that the newly developed phrase extraction model is better than noun phrases only, in terms of both intrinsic and extrinsic measures; phrase similarity learning appears to produce marginal improvement; and the community detection approach yields better phrase summaries with more accurate estimation of the number of student supporters.

In summary, the contribution of this work is threefold.

- We introduce a new phrase-based highlighting scheme for automatic summarization, a departure from prior work. It highlights the phrases in the human summary and also the semantically similar phrases in student responses. We create a new dataset annotated with this highlighting scheme¹.
- We push the boundary of a phrase-based summarization framework by using our highlighting scheme to enable identification of candidate phrases as well as estimation of phrase similarities with supervision, and by using community detection to group phrases into clusters.
- We conduct comprehensive evaluations in terms of both summary text quality, measured by ROUGE (Lin, 2004), and how well phrase summaries capture the most pressing student needs, measured by a new evaluation metric based on color matching.

¹This data set is publicly available at <http://www.coursemirror.com/download/dataset2>

2 Related Work

Work on automatic text summarization involves multiple granularities, ranging from keywords, phrases, to sentences. Traditional approaches have largely focused on sentence extraction (Martins and Smith, 2009; Berg-Kirkpatrick et al., 2011; Li et al., 2013) and document abstraction (Liu et al., 2015; Rush et al., 2015; Durrett et al., 2016; Nallapati et al., 2016). In both cases, the produced summary is expected to be cohesive and coherent. We deviate from this path and seek to directly generate a set of bullet points as a summary. Phrases are easy to search and browse like words but more meaningful, and fit better on the small screen of a mobile device compared to sentences (Ueda et al., 2000; Luo et al., 2015).

Our task setting differs from those of keyphrase extraction (Wu et al., 2005; Liu et al., 2009; Medelyan et al., 2009; Hasan and Ng, 2014; Kan, 2015). Of key importance is that each summary phrase is associated with a numerical value, indicating the number of students who raise the issue. This information is critical to course instructors for making informed choices. Intuitively our task setting bears similarity to word/phrase cloud (Yatani et al., 2011; Brooks et al., 2014), where the cloud gives greater prominence to words or phrases that appear frequently in the source text. The downside is that they do not take lexical variety into account or considering semantically-equivalent words/phrases.

A summarization system is expected to produce high quality summary phrases and accurate estimates of the number of student supporters for each phrase. Luo and Litman (2015) focus on extracting noun phrases from student responses, however there lacks a comprehensive evaluation of the results, taking the number of student supporters into account. Other related work on student responses includes collecting student responses using a mobile application named CourseMIRROR (Luo et al., 2015; Fan et al., 2015), determining the quality of a student reflective response and providing feedback (Luo and Litman, 2016), and extracting informative sentences from the student feedback (Luo et al., 2016).

Traditional approaches to summary annotation have been based on either sentence extracts or document abstracts (Loza et al., 2014; Xiong and Litman, 2014; Wang and Ling, 2016). An effective linkage between the document content and human summary on the micro level have been largely absent. Barker et al.(2016) partially address this challenge by linking a summary back to a group of sentences that support the summary. However, this linkage is weak since it tells only that there is one sentence or more supporting the summary within the group, without explicitly telling which one(s). Approaches such as Pyramid (Nenkova and Passonneau, 2004) have exploited creating Summary Content Units (SCUs) to establish such links and alleviate the challenge. The new highlighting scheme described in this work holds promise for establishing direct links between the phrases in student responses and those in the human summary, allowing us to develop a new evaluation metric based on color matching.

3 New Data and Annotation

When reviewing the student feedback, we observe that not all issues are equally important. Some teaching problems are more prominent than others. Summary phrases should naturally reflect the number of students who raise the issue. But until now a reasonable sized dataset has been missing for this type of summarization setting. In this work we create a new dataset for this purpose. This allows us to develop a class of summarization approaches that learn to extract summary phrases from the student responses and estimate the number of student supporters for each summary phrase.

Our dataset consists of two statistics courses offered in a research university for industrial engineers. After each lecture, the students were asked to respond to two carefully designed reflection prompts using a mobile application named CourseMIRROR²: 1) “Describe what you found most interesting in today’s class,” and 2) “Describe what was confusing or needed more detail.” For each course, two independent human annotators (native English speakers) with a statistics/mathematics background were recruited to create summaries for each lecture and prompt. The instructions we provide to the annotators include “*create a summary using 5 phrases and mark how many students semantically mentioned each phrase.*” We limit the number of summary phrases to 5 per lecture and prompt in order to provide a concise summary to the instructor. Note that the summary phrases are not limited to extracts; while abstracts

²<https://play.google.com/store/apps/details?id=edu.pitt.cs.mips.coursemirror>

and fusion of phrases are also possible, they are rare. We further ask the annotators to “highlight the corresponding phrases in the student responses which are semantically the same to the summary phrases using the same highlight colors.” The number of highlights in student responses should match the number of students who semantically mentioned the phrase. An example is illustrated in Table 1.

Note that L&L attempt to annotate the number of student supporters for summary phrases on a small dataset but without the highlighting scheme. We argue that the new highlighting scheme can provide many unique benefits. First, it allows us to track the “source phrases” that humans use to create the summary phrase. For example, the first summary phrase in Human Summary 1 of Table 1 is “central limit theorem.” It is created from a collection of phrases in the student responses, including “The central limit”, “central limit teorem” (a typo by the student), “CLT” (its abbreviation), and “Central Limit Thm” (another abbreviation). Naturally the highlighted source phrases lend themselves to a supervised approach to candidate phrase extraction. Second, the highlights inform us about the similarity and dissimilarity of phrases. For example, the source phrases that bear the same color are semantically similar to each other, whereas those with different colors are semantically dissimilar. In a similar vein, we develop a supervised approach that learns to predict the phrase similarity using highlights as guidance. Third, we are now able to accurately match the phrases in a system summary to those in a human summary, allowing the development of a novel summarization evaluation metric. For instance, assuming the system summary contains the phrase “Last problem about normalization” from S2 (Table 1), using the color highlights, we know that this phrase matches the human summary phrase “normalization (last example).” Such semantic matching between system and human summaries remains an elusive challenge for traditional summarization evaluation, but highlights make it an easy decision. Finally, the highlights on source texts indicate to what extent the information has been retained in the human summary. Specific to our task, we are interested to know the percentage of students whose responses are covered by the human summary. We define a student coverage score where a student is covered if and only if part of his/her response is highlighted. For example, in Table 1, S9 is considered not covered by Human Summary 1.

Basic statistics of the dataset are presented in Table 2.³ The student coverage scores (75.9% for Course A and 82.4% for Course B) highlight the effectiveness of the current annotation scheme, with a majority of students covered by the human summaries.

Course	# Students	# Lectures	Averaged by Lecture/Prompt				
			# Responses	# Words	Words Per Res.	# Highlights	Student Coverage
A	66	11	34.1	156.5	4.5	27.8	75.9%
B	74	24	41.9	161.8	3.7	37.2	82.4%

Table 2: Basic statistics of the dataset. Because the student responses and human summaries are created for each lecture and prompt, we take the average of the corresponding statistics.

4 Improved Phrase Summarization

So far we have motivated the need for a new dataset with a highlighting scheme for phrase-based summarization. We proceed by describing three improvements to the phrase-based summarization framework. Our first improvement involves a supervised approach to candidate phrase extraction (§4.1). Next, we learn to predict the pairwise phrase similarity (§4.2). Further, we explore a community detection algorithm to group the phrases into clusters (§4.3). We use the cluster size as an approximation to the number of student supporters for all the phrases within the cluster. L&L adopt LexRank (Erkan and Radev, 2004) to finally choose one representative phrase from each cluster. We follow the convention in this study. Note that our focus of this paper is not on developing new algorithms but to explore new capabilities that are enabled by the highlighting scheme. We thus perform direct comparisons with approaches described in L&L and leave comparisons to other approaches to future work. We present an intrinsic evaluation of each improvement in this section, followed by a comprehensive extrinsic evaluation in §5.

³While there are 22 lectures in total for Course A, unfortunately, only 11 of them have phrase highlighting.

4.1 Candidate Phrase Extraction

The phrase-based highlighting scheme lends itself to a supervised phrase extraction approach. In contrast, L&L used heuristics to extract noun phrases (NPs) only. This limitation has meant that informative non-NP phrases such as “how confidence intervals linked with previous topics” will be excluded from the summary, whereas uninformative NP phrases such as “the most interesting point” may be included.

We attempt to resolve this issue by formulating candidate phrase extraction as a word-level sequence labeling task. Concretely, we aim to assign a label to each word in the student responses. We choose to use the ‘BIO’ labeling scheme, where ‘B’ stands for the beginning of a phrase, ‘I’ for continuation of a phrase, ‘O’ for outside of a phrase. For example, “**The (B) central (I) limit (I)** and (O) **normal (B) approximations (I)**” illustrates the tagging of individual words, where the “The central limit” and “normal approximations” are two phrases highlighted by our annotators.

Local Features	<ul style="list-style-type: none"> • Word trigram within a 5-word window • Part-of-Speech tag trigram within a 5-word window • Chunk tag trigram within a 5-word window • Whether the word is in the prompt • Whether the word is a stopword • Label bigrams.
Global Features	<ul style="list-style-type: none"> • Total number of word occurrences (stemmed) • Rank of the word’s term frequency

Table 3: Local and global features for supervised phrase extraction. Local features are extracted within one student’s response. Global features are extracted using all student responses to a prompt in one lecture.

We choose to use the Conditional Random Fields (CRF) (Lafferty et al., 2001) as our sequence labeler⁴ and develop a number of features (Table 3) based on sentence syntactic structure and word importance to signal the likelihood of a word being included in the candidate phrase. During training, we merge the phrase highlights produced by two annotators in order to form a large pool of training instances. When two highlights overlap completely, e.g., “normal approximations” are marked by both annotators using different colors, we keep only one instance of the phrase, resulting in 1,115 and 2,682 instances for Course A and Course B respectively. When the highlights partially overlap, we use each phrase highlight as a separate training instance. In this and all the following experiments, we perform leave-one-lecture-out cross validation on all the lectures and report results averaged across folds. Table 4 presents the intrinsic evaluation results on the phrase extraction task. We calculate Precision (P), Recall (R) and F-measure (F) scores based on the exact match of system phrases to gold-standard phrases. While the sequence labeling approach and the features presented here are straightforward, they do produce a collection of candidate phrases with higher precision. It removes noun phrases that are commonly used by students but uninformative (e.g., “a little bit abstract”, “a problem with today’s topic”) as they were not highlighted by annotators. Phrase well-formedness is highly important to the summary quality, as evaluated in §5.

Candidate Phrase Extraction	Course A			Course B		
	P	R	F	P	R	F
L&L (NPs only)	0.426	0.633	0.503	0.538	0.714	0.609
Sequence Labeling with Highlights	0.692*	0.569*	0.618*	0.771*	0.743	0.753*

Table 4: Results of phrase extraction, intrinsically evaluated by comparing the system phrases to gold-standard phrases using exact match. The highest score in each column is shown in bold. * means the difference is significant with $p < 0.05$.

4.2 Similarity Learning

Accurately estimating pairwise phrase similarity plays an essential role in phrase-based summarization. Better similarity learning helps produce better phrase clusters, which in turn leads to more accurate estimation of the number of student supporters for each summary phrase. While a human annotator

⁴We use the implementation of Wapiti (Lavergne et al., 2010) with default parameters.

could distinguish the semantic similarity or dissimilarity of the phrase highlights, it remains unclear if a single similarity metric could fulfill this goal or if we may need an ensemble of different metrics.

L&L calculate the pairwise phrase similarity using SEMILAR (Rus et al., 2013) with the latent semantic analysis (LSA) trained on the Touchstone corpus (Ștefănescu et al., 2014). One drawback of this approach is that the similarity of phrases that do not appear in a background corpus cannot be captured. In this work we develop an ensemble of similarity metrics by feeding them into a supervised classification framework. We use the phrase highlights as supervision, where phrases of the same color are positive examples and those of different colors are negative examples. We experiment with a range of metrics for measuring lexical similarity, including lexical overlap (Rus et al., 2013), cosine similarity, LIN similarity (Miller, 1995), BLEU (Papineni et al., 2002), SimSum (Lin, 2004), Word Embedding (Goldberg and Levy, 2014), and LSA (Deerwester et al., 1990). LIN similarity is based on WordNet definitions. Lexical overlap, cosine similarity, BLEU, and SimSum are related to how many words the two phrases have in common, while Word Embedding and LSA both capture the phrase similarity in a low dimensional semantic space. Therefore, we use an ensemble of the above similarity metrics by feeding them as features in a SVM classification model, assuming it will be better suited for this task than the LSA alone. Table 5 presents the intrinsic evaluation results. LSA has a poor degree of coverage (low recall) with many phrase similarities not being picked up by the metric.

Pairwise Phrase Similarity	Course A			Course B		
	P	R	F	P	R	F
L&L (LSA)	0.904	0.665	0.730	0.878	0.506	0.584
Similarity Learning with Highlights	0.895	0.801*	0.833*	0.943*	0.768*	0.836*

Table 5: Results of predicting pairwise phrase similarity, measured using classification P/R/F.

4.3 Phrase Clustering

L&L use K-medoids for phrase clustering. It is a greedy iterative clustering algorithm (Kaufman and Rousseeuw, 1987), which may suffer from local minimal. We instead treat phrase clustering as a community detection problem. We define a **community** as a set of phrases that are semantically similar to each other, as compared to the rest of the phrases in student responses (Malliaros and Vazirgiannis, 2013). In our formulation, we consider each candidate phrase as a node in the network graph. We create an edge between two nodes if the two phrases are considered semantically similar to each other using the above similarity learning approach. Our goal is to identify tightly connected phrase communities in the network structure. The community size is used as a proxy for the number of students who semantically mention the phrase. Community detection has seen considerable success in tasks such as word sense disambiguation (Jurgens, 2011), medical query analysis (Campbell et al., 2014), and automatic summarization (Qazvinian and Radev, 2011; Mehdad et al., 2013).

Phrase Clustering	Course A	Course B
L&L (K-medoids)	82.2%	84.0%
Community Detection with OSLOM	85.2%*	88.8%*

Table 6: Results of phrase clustering measured by purity: ratio of number of phrases agreeing with the majority color in clusters.

We use OSLOM (Order Statistics Local Optimization Method, Lancichinetti et al., 2011) in this work. It is a widely used community detection algorithm that detects community structures (i.e., clusters of vertices) from a weighted, directed network. It optimizes locally the statistical significance of clusters with respect to a global null model during community expansion. We use an undirected version of OSLOM and set the p-value as 1.0 to encourage more communities to be identified⁵ since the number of vertices in the constructed graph is relatively small compared to large complex networks. The key feature of OSLOM is that it supports finding overlapped community structures and orphaned vertices, offering more flexibility in the clustering process than K-medoids. We want to investigate if the unique characteristics of OSLOM allow it to produce better phrase clusters, hence more accurate estimation of

⁵L&L set the number of clusters is to be the square root of the number of extracted phrases.

the number of student supporters. We conduct an intrinsic evaluation using purity, corresponding to the percentage of phrases in the cluster that agree with the majority color. Results are presented in Table 6. While this metric by itself is not thorough enough, it does highlight the strength of the community detection approach in generating cohesive clusters. One advantage of OSLOM we found is that it will treat a phrase different from any other phrase as a singleton, while this phrase must be assigned to one of the clusters in K-medoids, resulting in a noisy cluster.

5 Summary Evaluation

The previous section described three improvements to the phrase summarization framework. Next, we evaluate them on the end task of summarizing student course responses. The phrase summaries are evaluated along two dimensions: we expect ROUGE (Lin, 2004) to measure the informativeness of the summary text content (§5.1); we further propose a new metric to quantify to what extent the most pressing student needs have been captured in the summary (§5.2).

5.1 ROUGE

ROUGE measures the n-gram overlap between system and human summaries. In this work, we report R-1, R-2, and R-SU4 scores, which respectively measure the overlap of unigrams, bigrams, and unigrams plus skip bigrams with a maximum distance of 4 words. These are metrics commonly used in the DUC and TAC competitions (Dang and Owczarzak, 2008). We implement the phrase summarization framework described in (Luo and Litman, 2015), named as **PhraseSum**. Further, we include **LexRank** (Erkan and Radev, 2004) as a competitive baseline. LexRank is a graph-based summarization approach based on eigenvector centrality. It has demonstrated highly competitive performance against the **PhraseSum** on a prior dataset (Luo and Litman, 2015). The summary is limited to 5 phrases or less in all experiments. Note that, the summary length is set independently of the number of clusters. If the number of clusters produced in §4.3 is less than 5, the phrase number is equal to the cluster number.

Course	System	R-1			R-2			R-SU4		
		P	R	F	P	R	F	P	R	F
A	LexRank	.276*	.511	.348*	.118*	.245	.154	.077*	.260	.106*
	PhraseSum	.402	.466	.415	.170	.208	.178	.162	.222	.160
	SequenceSum	.600*	.448	.493*	.307*	.231	.249*	.368*	.225	.244*
	SimSum	.597*	.460	.504*	.302*	.241	.260*	.355*	.227	.249*
	CDSum	.634*	.435	.499*	.335*	.229	.262*	.404*	.210	.250*
B	LexRank	.357*	.560	.429*	.187*	.304*	.227	.129*	.290	.168*
	PhraseSum	.492	.545	.508	.231	.258	.239	.234	.283	.241
	SequenceSum	.618*	.485*	.531	.347*	.267	.294*	.385*	.238*	.274
	SimSum	.618*	.500*	.543	.353*	.284	.309*	.379*	.250	.285*
	CDSum	.702*†	.480*	.550*	.433*†	.279	.324*	.500*†	.240*	.293*

Table 7: Summarization Performance. **SequenceSum** means replacing the syntax phrase extraction in the PhraseSum baseline with the supervised sequence labeling phrase extraction. **SimSum** means replacing not only the phrase extraction but also the similarity scores using the supervised models. **CDSum** means using all three proposed techniques including the community detection. * indicates that the difference is statistically significant compared to PhraseSum with $p < 0.05$. † means that the improvement over SequenceSum is statistically significant with $p < 0.05$.

The summarization performance is shown in Table 7 (the caption explains the system names). The PhraseSum baseline, compared to LexRank, gets better P and F scores for all three ROUGE metrics for both courses, and the improvement of P is significant. This is the same as the findings in (Luo and Litman, 2015), and verifies our implementation of their model. For our enhancements of PhraseSum, the proposed supervised phrase extraction (SequenceSum) significantly improves P and thus improves (mostly significantly) F as well. SimSum is slightly better than SequenceSum for R and F, however, it is not significant using a two-tailed paired t-test. It suggests that a supervised method is not necessarily better than an unsupervised model in terms of the end-task performance, and its improvement over the

PhraseSum baseline is mainly due to the supervised phrase extraction step. In fact, the predicted similarity scores using the similarity learning model and the LSA model are highly correlated to each other ($r = 0.852$, $p < 0.01$) although it has a better classification performance (Table 5). Although CDSum is not significantly different from SequenceSum for the Course A, it does improve P significantly for all three ROUGE metrics for Course B. One possible explanation is that the latter course has a larger number of student responses, and thus benefits more from the community detection as the graph is larger.

5.2 A New Metric based on Color Matching

Our goal is to create a comprehensive evaluation metric that takes into account the following two factors.

- **Phrase matching.** While ROUGE is a classic summarization evaluation metric, it trivially compares the system vs. human summaries based on surface text form. In contrast, the phrase highlights allow us to accurately match the phrases in the system summary to those in the human summary based on color matching. This is due to two facts: first, our methods are extractive-based and all candidate phrases are extracted from the student responses; second, in the new highlighting scheme, the annotators are asked to highlight both the human summary phrase and any phrases in the student responses that are semantically the same with the summary phrase using the same color. It thus becomes easy to track the colors of the extracted phrases and verify if they match any of those in the human summary.
- **Student supporters.** Each summary phrase is tagged with the number of students who raise the issue. For human summary, this number is created by human annotators. For system summary, we approximate this number using the size of the cluster, from which the summary phrase is extracted.

Our proposed new metric resembles precision, recall, and F-measure. We define the true positive (TP) as the number of *shared colors* between system and human summaries. Each color is weighted by the number of student supporters, taken as the smaller value between system and human estimates. The *precision* is defined as TP over the total number of colors in the *system* summary, each weighted by system estimates; while *recall* is defined as TP over the total number of colors in the *human* summary, each weighted by human estimates. For example, assuming the phrases in the human summary are colored and tagged with estimates on student support: yellow/12, green/9, red/6, blue/5, magenta/3; similarly the phrases in the system summary are colored and tagged: yellow/11+3, green/17, red/7, blue/7. There are two phrases in the system summary that bear the same color, we thus add up the system estimates into yellow/11+3 (see Human Summary 1 in Table 1 and SequenceSum in Table 9). There are 4 shared colors between system and human summaries. The true positive is calculated as: $12 + 9 + 6 + 5 = 32$. The precision is $32 / ((11 + 3) + 17 + 7 + 7) = 0.711$, and recall is $32 / (12 + 9 + 6 + 5 + 3) = 0.914$. The F-measure is calculated as the harmonic mean of precision and recall scores.

The performance is shown in Table 8. Similar to the ROUGE evaluation, SequenceSum improves the P and F significantly. Now, CDSum not only significantly improves P, but also F for Course B.

	Course A			Course B		
	P	R	F	P	R	F
PhraseSum	.349	.615	.437	.485	.747	.576
SequenceSum	.626*	.642	.614*	.698*	.757	.717*
SimSum	.602*	.636	.595*	.711*	.753	.723*
CDSum	.643*	.634	.613*	.777*†	.762	.759*†

Table 8: Evaluation based on the new metric of color matching. P, R, and F are averaged by the annotators.

5.3 Example Summaries

The automatic summaries generated by different systems for the same example in Table 1 are shown in Table 9. The PhraseSum baseline extracts unnecessary content, which could be eliminated by the supervised phrase extraction model. For example, including “the example after” before “central limit theorem” makes it too specific. The “collapse” effect with a large cluster with unrelated items (Basu et al., 2013) can also be illustrated (e.g., the quantitative numbers for the phrase “i” in PhraseSum and

“q-q plot” in “SequenceSum” are much larger than the gold standard). This is solved by the community detection algorithm where such bigger clusters will not be considered as a single community.

PhraseSum	SequenceSum	CDSum
- i [40]	- q-q plot ^g [17]	- central limit theorem ^y [11]
- the example after central limit theorem ^y [12]	- central limit theorem ^y [11]	- q-q plot ^g [10]
- q q plot ^g [9]	- normal approximation to binomial ^b [7]	- sampling distributions ^r [7]
- the fact that we can sample as many as we want [9]	- sampling distributions ^r [7]	- normal approximation to binomial ^b [5]
- last problem about normalization ^m [6]	- clt ^y [3]	- nothing [4]

Table 9: Example system summaries for the example in Table 1. Note, the highlights in these summaries are NOT annotated by human after they are generated. Instead, they are automatically extracted from the dataset (§5.2).

6 Conclusion and Future Work

In this work, we introduce a new phrase-based highlighting scheme for automatic summarization. It highlights the phrases in the human summary and also the corresponding phrases in student responses. Enabled by the highlighting scheme, we improved the phrase-based summarization framework proposed by Luo and Litman (2015) by developing a supervised candidate phrase extraction, learning to estimate the phrase similarities, and experimenting with different clustering algorithms to group phrases into clusters. We further introduced a new metric that offers a promising direction for making progress on developing automatic summarization evaluation metrics. Experimental results show that our proposed methods not only yield better summarization performance evaluated using ROUGE, but also produce summaries that capture the pressing student needs. Future work includes thorough comparison with other approaches and extending the current research to multiple courses and other summary lengths in order to test the generalizability. We also plan to supplement our ROUGE scores with human evaluations of system summaries.

Acknowledgements

This research is supported by an internal grant from the Learning Research and Development Center at the University of Pittsburgh. We thank Jingtao Wang and Xiangmin Fan for developing the CourseMIRROR mobile system. We thank Fan Zhang and Huy Nguyen for valuable suggestions about the proposed summarization algorithm. We also thank anonymous reviewers for insightful comments and suggestions.

References

- Emma Barker, Monica Lestari Paramita, Ahmet Aker, Emina Kurtic, Mark Hepple, and Robert Gaizauskas. 2016. The sensei annotated corpus: Human summaries of reader comment conversations in on-line news. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 42–52, Los Angeles, September. Association for Computational Linguistics.
- Sumit Basu, Chuck Jacobs, and Lucy Vanderwende. 2013. Powergrading: a clustering approach to amplify human effort for short answer grading. *Transactions of the Association for Computational Linguistics*, 1:391–402.
- Taylor Berg-Kirkpatrick, Dan Gillick, and Dan Klein. 2011. Jointly learning to extract and compress. In *Proceedings of ACL*, pages 481–490, Portland, Oregon, USA.
- David Boud, Rosemary Keogh, David Walker, et al. 2013. *Reflection: Turning experience into learning*. Routledge.
- Bill J Brooks, Debra M Gilbuena, Stephen Krause, and Milo D Koretsky. 2014. Using word clouds for fast, formative assessment of students’ short written responses. *Chemical Engineering Education*, 48(4):190–198.
- William Campbell, Elisabeth Baseman, and Kara Greenfield. 2014. Content+context=classification: Examining the roles of social interactions and linguist content in Twitter user classification. In *Proceedings of the Second Workshop on Natural Language Processing for Social Media*, pages 59–65, Dublin, Ireland.

- Dan Ștefănescu, Rajendra Banjade, and Vasile Rus. 2014. Latent semantic analysis models on wikipedia and tasa. In *Proceedings of LREC*, pages 26–31, Reykjavik, Iceland.
- Hoa Trang Dang and Karolina Owczarzak. 2008. Overview of the TAC 2008 update summarization task. In *Proceedings of TAC*, pages 1–16.
- Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391.
- Greg Durrett, Taylor Berg-Kirkpatrick, and Dan Klein. 2016. Learning-based single-document summarization with compression and anaphoricity constraints. In *Proceedings of ACL*, pages 1998–2008, Berlin, Germany.
- Günes Erkan and Dragomir R. Radev. 2004. LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22(1):457–479.
- Xiangmin Fan, Wencan Luo, Muhsin Menekse, Diane Litman, and Jingtao Wang. 2015. CourseMIRROR: Enhancing large classroom instructor-student interactions via mobile interfaces and natural language processing. In *Works-In-Progress of ACM Conference on Human Factors in Computing Systems*. ACM.
- Yoav Goldberg and Omer Levy. 2014. word2vec explained: Deriving Mikolov et al.’s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722*.
- Kazi Saidul Hasan and Vincent Ng. 2014. Automatic keyphrase extraction: A survey of the state of the art. In *Proceedings of ACL*, pages 1262–1273, Baltimore, Maryland.
- David Jurgens. 2011. Word sense induction by community detection. In *Proceedings of TextGraphs-6 Workshop*, pages 24–28, Portland, Oregon.
- Min-Yen Kan. 2015. Keywords, phrases, clauses and sentences: topicality, indicativeness and informativeness at scales. In *Proceedings of the ACL 2015 Workshop on Novel Computational Approaches to Keyphrase Extraction*, page 1, Beijing, China.
- Leonard Kaufman and Peter Rousseeuw. 1987. Clustering by means of medoids. *Statistical Data Analysis Based on the L1-Norm and Related Method*, pages 405–416.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, pages 282–289, San Francisco, CA, USA.
- Andrea Lancichinetti, Filippo Radicchi, José J Ramasco, and Santo Fortunato. 2011. Finding statistically significant communities in networks. *PloS one*, 6(4):e18961.
- Thomas Lavergne, Olivier Cappé, and François Yvon. 2010. Practical very large scale CRFs. In *Proceedings the 48th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 504–513. Association for Computational Linguistics, July.
- Chen Li, Fei Liu, Fuliang Weng, and Yang Liu. 2013. Document summarization via guided sentence compression. In *Proceedings of EMNLP*, pages 490–500, Seattle, Washington, USA.
- Chin-Yew Lin. 2004. ROUGE: a package for automatic evaluation of summaries. In *Proceedings of the Workshop on Text Summarization Branches Out*, volume 8. Barcelona, Spain.
- Zhiyuan Liu, Peng Li, Yabin Zheng, and Maosong Sun. 2009. Clustering to find exemplar terms for keyphrase extraction. In *Proceedings of EMNLP*, pages 257–266, Stroudsburg, PA, USA.
- Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward abstractive summarization using semantic representations. In *Proceedings of NAACL*, pages 1077–1086, Denver, Colorado.
- Vanessa Loza, Shibamouli Lahiri, Rada Mihalcea, and Po-Hsiang Lai. 2014. Building a dataset for summarization and keyword extraction from emails. In *Proceedings of LREC*, pages 2441–2446, Reykjavik, Iceland.
- Wencan Luo and Diane Litman. 2015. Summarizing student responses to reflection prompts. In *Proceedings of EMNLP*, pages 1955–1960, Lisbon, Portugal.
- Wencan Luo and Diane Litman. 2016. Determining the quality of a student reflective response. In *Proceedings 29th International FLAIRS Conference*, Key Largo, FL.

- Wencan Luo, Xiangmin Fan, Muhsin Menekse, Jingtao Wang, and Diane Litman. 2015. Enhancing instructor-student and student-student interactions with mobile interfaces and summarization. In *Proceedings of NAACL: Demonstrations*, pages 16–20, Denver, Colorado.
- Wencan Luo, Fei Liu, Zitao Liu, and Diane Litman. 2016. Automatic summarization of student course feedback. In *Proceedings of NAACL*, pages 80–85, San Diego, California.
- Fragkiskos D. Malliaros and Michalis Vazirgiannis. 2013. Clustering and community detection in directed networks: A survey. *CoRR*, abs/1308.0971.
- Andre Martins and Noah A. Smith. 2009. Summarization with a joint model for sentence extraction and compression. In *Proceedings of the Workshop on Integer Linear Programming for NLP*, pages 1–9, Boulder, Colorado.
- Olena Medelyan, Eibe Frank, and Ian H. Witten. 2009. Human-competitive tagging using automatic keyphrase extraction. In *Proceedings of EMNLP*, pages 1318–1327, Stroudsburg, PA, USA.
- Yashar Mehdad, Giuseppe Carenini, Frank Tompa, and Raymond T. NG. 2013. Abstractive meeting summarization with entailment and fusion. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 136–146, Sofia, Bulgaria.
- Muhsin Menekse, Glenda Stump, Stephen J. Krause, and Michelene T.H. Chi. 2011. The effectiveness of students daily reflections on learning in engineering context. In *Proceedings of the American Society for Engineering Education Annual Conference*, Vancouver, Canada.
- George A Miller. 1995. WordNet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Frederick Mosteller. 1989. The ‘muddiest point in the lecture’ as a feedback device. *On Teaching and Learning: The Journal of the Harvard-Danforth Center*, 3:10–21.
- Ramesh Nallapati, Bing Xiang, and Bowen Zhou. 2016. Sequence-to-sequence rnns for text summarization. *CoRR*, abs/1602.06023.
- Ani Nenkova and Rebecca Passonneau. 2004. Evaluating content selection in summarization: The pyramid method. In Daniel Marcu Susan Dumais and Salim Roukos, editors, *Proceedings of NAACL*, pages 145–152, Boston, Massachusetts, USA.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BIEU: a method for automatic evaluation of machine translation. In *Proceedings of ACL*, pages 311–318, Philadelphia, Pennsylvania, USA.
- Vahed Qazvinian and Dragomir R. Radev. 2011. Learning from collective human behavior to introduce diversity in lexical choice. In *Proceedings of ACL*, pages 1098–1108, Portland, Oregon, USA.
- Vasile Rus, Mihai Lintean, Rajendra Banjade, Nobal Niraula, and Dan Stefanescu. 2013. Semilar: The semantic similarity toolkit. In *Proceedings of ACL: System Demonstrations*, pages 163–168, Sofia, Bulgaria.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of EMNLP*, pages 379–389, Lisbon, Portugal.
- Yoshihiro Ueda, Mamiko Oka, Takahiro Koyama, and Tadanobu Miyauchi. 2000. Toward the “at-a-glance” summary: Phrase-representation summarization method. In *Proceedings of COLING*, pages 878–884, Stroudsburg, PA, USA.
- Gerard Van den Boom, Fred Paas, Jeroen JG Van Merriënboer, and Tamara Van Gog. 2004. Reflection prompts and tutor feedback in a web-based learning environment: effects on students’ self-regulated learning competence. *Computers in Human Behavior*, 20(4):551 – 567.
- Lu Wang and Wang Ling. 2016. Neural network-based abstract generation for opinions and arguments. In *Proceedings of NAACL*, pages 47–57, San Diego, California.
- Yi-fang Brook Wu, Quanzhi Li, Razvan Stefan Bot, and Xin Chen. 2005. Domain-specific keyphrase extraction. In *Proceedings of CIKM*, pages 283–284, New York, NY, USA.
- Wenting Xiong and Diane Litman. 2014. Empirical analysis of exploiting review helpfulness for extractive summarization of online reviews. In *Proceedings of COLING*, pages 1985–1995, Dublin, Ireland.
- Koji Yatani, Michael Novati, Andrew Trusty, and Khai N. Truong. 2011. Review Spotlight: A user interface for summarizing user-generated reviews using adjective-noun word pairs. In *Proceedings of CHI*, pages 1541–1550, New York, NY, USA.