

Overview of the 2021 Key Point Analysis Shared Task

Roni Friedman*, Lena Dankin*, Yufang Hou, Ranit Aharonov,
Yoav Katz and Noam Slonim

IBM Research

{roni.friedman-melamed, lenad, ranit.aharonov2, katz, noams}@il.ibm.com
yhou@ie.ibm.com

Abstract

We describe the 2021 Key Point Analysis (KPA-2021) shared task on key point analysis that we organized as a part of the 8th Workshop on Argument Mining (ArgMining 2021) at EMNLP 2021. We outline various approaches and discuss the results of the shared task. We expect the task and the findings reported in this paper to be relevant for researchers working on text summarization and argument mining.

1 Introduction

Key Point Analysis (KPA) was introduced in Bar-Haim et al. (2020a,b) as a challenging NLP task with tight relations to Computational Argumentation, Opinion Analysis, and Summarization, and with many practical applications (Bar-Haim et al., 2021). Given a potentially large collection of relatively short, opinionated texts focused on a topic of interest, the goal of KPA is to produce a succinct list of the most prominent key-points (KPs) in the input corpus, along with their relative prevalence. Thus, the output of KPA is a bullet-like summary, with an important quantitative angle. Successful solutions to KPA can be used to gain better insights from public opinions as expressed in social media, surveys, and so forth, giving rise to a new form of a communication channel between decision makers and people that might be impacted by the decision.

Various requirements govern the value of the KPA output. KPs are expected to be succinct, non-redundant, capturing points that are central to the topic of interest, and reflecting a clear stance towards that topic. Ideally, they should be at the right granularity for summarising the input data – not too specific and yet still informative and not overly general. In addition, accurate mapping of input texts to KPs is obviously essential. First, to ensure a reliable estimate of the prevalence of each key

point. And second, to enable the user to drill-down, to gain a deeper understanding of the issues underlying each key point, as expressed by the input texts mapped to that key point.

The goal of the KPA-2021 shared task was to further increase the attention of the NLP community to this emerging task, while enriching the space of existing KPA solutions. Since providing a complete KPA solution is challenging, we divided the task into two tracks, enabling teams to participate only in the first, relatively simpler track. Specifically, in the first track, referred to as *Matching Track*, KPs are given as part of the input, and the task is focused on mapping input text to these KPs. In contrast, in the second track, referred to as *Generation Track*, no KPs are provided and the task requires to further generate the KPs. The data being considered for evaluation on both tracks are crowd sourced arguments on three debatable topics.

The nature of the KPA task requires to consider various evaluation measures to estimate the quality of the results across different dimensions. In this paper we report and analyze the results of 22 models submitted by 17 teams to the KPA-2021 shared task across these different evaluation measures. We also discuss lessons learned from these multiple submissions, that can guide future development of new KPA algorithms and associated evaluation methods.

2 Related Work

Our shared task is based on Bar-Haim et al. (2020a,b), which introduced the problem of KPA and proposed an extractive approach to address the task. In particular, we focus on benchmarking two sub tasks: (1) Key point matching; and (2) Key point generation and matching.

For the matching task Bar-Haim et al. (2020b) uses RoBERTa (Liu et al., 2019) based classifier, trained on argument-KP pairs. For the generation task (Bar-Haim et al., 2020b) employs the follow-

*These authors equally contributed to this work.

ing steps: First, compiling a list of high quality arguments that serve as KP candidates. Next, a match confidence score is calculated between all KP candidate-argument pairs, as well as between all the KP candidates to themselves. Finally, several filtering and ordering strategies are employed to maintain the most salient, yet non-overlapping KP candidates to serve as KPs.

Below we briefly review the work closely related to KPA in a few fields.

Multi-document summarization (MDS). Most prior work on MDS has been focused on the news, Wikipedia, and scientific literature domains (Dang, 2005; Baumel et al., 2016; Liu et al., 2018; Lu et al., 2020). MDS techniques are typically query-based and extractive, where models select a set of non-redundant sentences that are most relevant to the query to include in the summary (Xu and Lapata, 2020). A few studies (Liu et al., 2018; Shapira et al., 2021) have explored generating the most salient queries the users might be interested in. However, it is still unclear how these salient queries/points are represented in the source documents. KPA addresses this problem by reporting the prevalence of each KP in the input data.

Opinion and argument summarization. Recently there has been an increasing interest in summarizing opinions expressed in various reviews (Amplayo and Lapata, 2021; Elsahar et al., 2021) or argumentative text (Wang and Ling, 2016; Chen et al., 2019; Syed et al., 2020). KPA contributes to this line of work by adding a quantitative dimension which reflects the distribution of opinions in the examined data. For a detailed discussion of the relation between KPA and argument clustering and summarization see (Bar-Haim et al., 2020a,b).

3 Task Details

3.1 Task Description

Given a collection of arguments, with either pro or con stance towards a certain debatable topic, the goal of the KPA-2021 task is to generate a KP-based quantitative summary. The *Generation Track* requires models to perform the entire task – i.e., generate a list of KPs for each topic and stance, and then predict for each argument in this topic and stance the confidence by which it is matched to each KP. Since this is rather challenging we also established the *Matching Track*, in which a set of KPs, generated by an expert per topic and stance,

is given as part of the input, requiring models to only predict a match confidence score for each argument in the input, towards each of the provided KPs. In both tracks, the confidence score expresses the model’s confidence that the KP represents the essence of the argument, and can be associated with it for the purpose of a KP-based quantitative summary. Since in the ground truth data (defined in the next section), only 5% of the arguments map to multiple KPs, we take a simplifying approach in the evaluation, where we consider for each argument only the KP with the highest confidence score as true positive.

Since in the ground truth data, a significant portion of the arguments are not matched to any KP, it is not enough for a model to rank the KPs per argument. Rather, it should also provide a threshold on the confidence that determines whether the highest scoring KP is matched to the argument, or whether this argument has no matching KP. In this shared task, we opted for a simpler approach, in which the submitted models are not required to determine the threshold. Rather, we evaluate all models by considering only the top 50% of the argument-KP pairs, as ranked by the confidence scores of all pairs.

The full terms of the KPA-2021 task can be found online¹.

3.2 The ArgKP-2021 data set

For KPA-2021 we use a data set covering a set of debatable topics, where for each topic and stance, a set of triplets of the form <argument, KP, label> is provided. The data set is based on the ArgKP data set (Bar-Haim et al., 2020a), which contains arguments contributed by the crowd on 28 debatable topics, split by their stance towards the topic, and KPs written by an expert for those topics. As described in (Bar-Haim et al., 2020a), crowd annotations were collected to determine whether a KP represents an argument, i.e., is a match for an argument, and based on these annotations every argument-KP pair received a label that can take three values: (i) positive - 60% or more of the annotations determined there is a match; (ii) negative - 15% or less of the annotations determined there is a match; and (iii) undecided - neither positive nor negative. For example, in the pro stance towards the topic "Homeschooling should be banned", and

¹https://github.com/ibm/KPA_2021_shared_task

for the KP "Mainstream schools are essential to develop social skills", the argument "children can not learn to interact with their peers when taught at home" was labeled as positive and the argument "homeschooling causes a decline in education" as negative. The arguments in ArgKP are a subset of the IBM-ArgQ-Rank-30kArgs data set (Gretz et al., 2020), which consists of 30K crowd-sourced arguments on 71 controversial topics, where each argument is labeled for its stance towards the topic as well as for its quality.

The ArgKP data set, which is publicly available, served as the training data for our task. The 28 topics in the data set were split to train (24 topics) and dev (4 topics) sets and were offered to the participants with a recommendation to use as such in the preparation of their models. For *Generation Track*, in addition to the ArgKP data set, the participants were encouraged to use the remaining arguments from the IBM-ArgQ-Rank-30kArgs data set.

For a test set, we extended ArgKP, adding three new debatable topics, that were also not part of IBM-ArgQ-Rank-30kArgs. The test set was collected specifically for KPA-2021, and was carefully designed to be similar in various aspects to the training data². For each topic, crowd sourced arguments were collected, expert KPs generated, and match/no match annotations for argument/KP pairs obtained, resulting in a data set compatible with the ArgKP format. Arguments collection strictly adhered to the guidelines, quality measures, and post processing used for the collection of arguments in IBM-ArgQ-Rank-30kArgs (Gretz et al., 2020), while the generation of expert KPs, collection of match annotations, and final data set creation strictly adhered to the manner in which ArgKP was created (Bar-Haim et al., 2020a).

The full data set, i.e., ArgKP together with the test set, is termed ArgKP-2021 and is available online³. See Table 1 for data statistics.

3.3 Evaluation Metrics

3.3.1 Matching Track

For the submitted prediction files, we extract argument-KP pairs as follows. First, each argument is paired with the highest scoring KP as-

²Those aspects include, among others – number of arguments in each topic; the distribution of their quality and length; the amount and style of expert KPs produced for them; as well the coverage of those as obtained by crowd labeling.

³https://www.research.ibm.com/haifa/dept/vst/debating_data.shtml

signed to it (randomly chosen in case of a tie)⁴. Second, the top 50% of the argument-KP pairs are extracted for each topic and stance as an ordered list, based on the confidence score. Finally, given the ground truth data, Average Precision (AP, Turpin and Scholer (2006)), is calculated per topic and stance. Note that a perfect score of 1 can only be achieved when all top 50% matches are correct. We chose AP as the evaluation metric, since it supports evaluating the correlation between a model’s confidence and prediction success. We obtain the mean AP (mAP) by macro averaging over all topics and stances so that each topic and stance has the same effect over final score.

When evaluating the models, we need to address the issue of the undecided argument-KP pairs in the labeled data. Considering a prediction of undecided pairs as a false prediction is too strict, since > 15% of the annotators found this match correct, and hence it is not entirely wrong. Discarding predictions of undecided pairs from the set of a model’s predictions, as was done in (Bar-Haim et al., 2020a), is not relevant for KPA-2021, as it would have rendered the evaluation process inconsistent across different models. We therefore relied on two measures for this track - **relaxed** and **strict** mAP. In the strict evaluation score, undecided pairs are considered as wrong predictions, while in the relaxed evaluation score they are considered as correct, i.e., true positives. The final rank of a model in this track is the average between the rank of the model in the strict score and its rank in the relaxed score.

3.3.2 Generation Track

Models submitted to this track were evaluated on two aspects: The first is how well the set of generated KPs serves as a summary of possible arguments for this topic and stance. The second is how well does the model match between the arguments and the generated KPs. We note that evaluating only how well the arguments match the generated KPs is insufficient since one may obtain relatively high mAP scores by generating redundant and/or very general and non-informative KPs. It is therefore crucial to establish the quality of the set of KPs as a standalone evaluation. Due to the nature of this track, no ground truth data is available, and we therefore perform the evaluation of the submit-

⁴In such cases, each experiment was repeated 10 times with different random seeds, and we report the resulting average. Standard deviation for those rare cases is insignificant.

Split	#Args	#KPs	Positive	Negative	Undecided	Strict coverage	Relaxed coverage
train (24)	5583	207	0.17	0.67	0.16	0.72	0.94
dev (4)	932	36	0.18	0.65	0.18	0.72	0.95
test (3)	723	33	0.14	0.73	0.13	0.69	0.91

Table 1: The ArgKP-2021 data set. The number of topics in each split is given in brackets. Coverage represents the fraction of arguments that have at least one matching KP. The strict view considers only argument-KP pairs with a positive label, while the relaxed view considers all pairs with a positive or undecided label. See text for details.

ted results using crowd annotations via the Appen platform⁵. The annotations were generated by a selected group of annotators which performed well in previous tasks, and further obtained good results in sanity-test questions and demonstrated high agreement with other annotators (Gretz et al. (2020), Toledo et al. (2019)).

The first evaluation considers the model’s ability to correctly match arguments to KPs. As in *Matching Track*, an argument is matched to a single KP, and only the top 50% matches are considered for evaluation. Unlike *Matching Track*, here the evaluation of the match must be retrospective, since the KPs are provided by the submitted models. Following Bar-Haim et al. (2020b), we presented each argument to the annotators with its top matching KP, and asked if the KP matches the argument. Possible answers are ‘Yes’/‘No’/‘Faulty KP’. We consider ‘Faulty KP’ as no match. Since here we directly ask regarding the predicted match, rather than showing annotators all possible pairs (as was done when generating the ground truth data for the *Matching Track*), we generate a binary label from the annotations by the majority vote, and compute a single mAP score.

The second evaluation aspect concerns the quality of the set of KPs generated by the submitted model. Here we asked annotators to perform a pairwise comparison between two sets of KPs, each taken from a different model, on the same topic and stance. We perform such pairwise comparisons between all pairs of top models according to *Generation Track* mAP score. Each KP set is sorted, such that the more prevalent KPs according to this model’s matching are higher in the list. KPs that match 1% or less of the arguments are not presented to the annotators.

The annotators are first asked to determine for each KP set what is the stance of the set towards the topic (All support the topic/All contest the topic/Most support the topic/Most contest the

topic). This initial question is not positioned as a comparative question. However, annotators’ answers are compared to the stance of the arguments of the KP set, and in case one model is correct and the other is not (either all or some KPs are judged not to be in the correct stance), this is counted as a win for the former. Two comparative questions follow: *Which of the two summaries is less redundant?* and *Which of the two summaries captures points that are more central to the topic?*, where the latter question aims for the relevance of the KP set. Possible answers for those questions are *Summary A/Summary B/Both are equal in this respect*. For each of the three questions, ranking is calculated using the Bradley-Terry model (Bradley and Terry, 1952), which predicts the probability of a given participant to win a paired comparison, based on previous paired comparison results of multiple participants, and thus allows ranking them. The final rank of a model in the KP set quality aspect is the average of the three ranks.

Finally, a model’s rank in the *Generation Track* is the average of its rank in the mAP evaluation aspect and its rank in the evaluation of the KP set quality. In case of a tie between the two ranks, mAP score is given preference.

4 Submitted models

In total, 17 models were submitted to *Matching Track*, and 5 models to *Generation Track*. Note that some teams did not provide description for their model. Still, we show their results in the results section, but we cannot provide their descriptions.

4.1 Matching Track

As in other NLP related tasks, transformer based approaches are preferable among the participants, with a clear dominance of RoBERTa (Liu and Lapata, 2019). None of the participants found the undecided pairs as useful for training, thus training was focused only on positive and negative pairs. Most of the participants used the topic as an in-

⁵www.appen.com.

put to their models, and ignored the stance. We also note that some of the models (SMATCHToPR, MatchTStm) utilized contrastive learning, taking advantage of the fact that arguments are provided with matched and unmatched KPs. **SiamBERT**, though not directly by contrastive learning, also shows two key points for one argument in training.

SMATCHToPageRank (hereafter as **SMATCHToPR**) (Alshomary et al., 2021) aims to learn an embedding space where matching argument-KP pairs are closer than non-matching pairs. This model has two inputs: (1) the argument; and (2) the concatenation of the KP and the topic. The tokens of each input are embedded using RoBERTa-large. Next, those inputs are passed to a siamese Sentence-BERT network that is trained with a contrastive loss.

Team **NLP@UIT** use ALBERT XXLARGE (Lan et al., 2019) pre-trained model for sentence-pair classification for pairs of argument and KP. Five models are fine-tuned on different cross-validation splits, and for inference the average of all models is considered as the final score.

Modern Talking - RoBERTa-Base (hereafter as **ModrnTalk**) (Reimer et al., 2021) is a RoBERTa-Base model fine-tuned on the concatenation of arguments-kps pairs. At training, the model is warmed up on 6% of the data.

Enigma (Kapadnis et al., 2021) is based on a DeBERTa-Large (He et al., 2020) pre-trained model. The input to the model is the concatenation of KP, argument, and topic. The output of the DeBERTa model is concatenated with the POS tags representation of the three text pieces, and then fed to two more dense layers.

Matching the statement (hereafter as **MatchT-Stm**) (Phan et al., 2021) embeds each argument and KP separately along with their context that is represented by the topic and stance. RoBERTa is utilized to extract all text embeddings. Next, the statement pairs are passed to a fully connected network. Finally, a contrastive learning is employed by grouping matching and non matching arguments for the same KP in a single batch.

poledro (Cosenza, 2021) is a fine-tuned ALBERT-XXLARGE model, similarly to NLP@UIT. The input to the model is a concatenated argument-KP pair.

XLNet is a vanilla XLNet-Large auto-regressive model (Yang et al., 2019) fine-tuned on argument-KP pairs.

RoBERTa-large-concat (hereafter as **RoBERTaLC**) is a fine-tuned RoBERTa-large model. The input to the model is the concatenation of the topic, the KP, and the argument. Training data is balanced between positive and negative match examples by over sampling positive instances.

Siamese BERT classifier (hereafter as **SiamBERT**) is a siamese BERT (Devlin et al., 2019) based model trained on pairs of arguments and KPs. At each training step, the same argument with two different KPs, one that matches and one that doesn't, is passed to the model, and it learns to prefer the matching pair.

4.2 Generation Track

SMATCHToPR applies an extractive approach. First, arguments with high quality scores (using (Toledo et al., 2019; Gretz et al., 2020)) are selected as KP candidates. Next, arguments with high matching score (using *Matching Track* models) are connected to form an undirected graph, where the nodes are the arguments. Finally, an importance score is calculated for each node, and most important nodes are selected as KPs. Note that KP candidates that have high confidence scores with other candidate are ignored, to ensure diversity.

Enigma approaches this task as an abstractive summarization task. First, Pegasus (Zhang et al., 2020) is fine-tuned with an input of concatenated arguments and topics. The set of corresponding KPs are provided as a ground truth summary. For each pair of arguments and the corresponding topic, several summaries were generated. Next, those summaries are ranked by the ROUGE-1 score with respect to the input pair, and the top 5 summaries are selected as KPs for each stance.

XLNet applies the XLNet model developed for *Matching Track* on all possible pairs of arguments, so that each argument is evaluated as a potential KP summarizing another argument. Arguments with the highest average confidence scores are selected as KPs.

5 Results

In addition to the submitted models, we evaluate the state of the art RoBERTa based model described in (Bar-Haim et al., 2020b) and in section 2, termed hereafter BarH.

5.1 Matching Track

Table 2 describes the evaluation results for the *Matching Track*. By definition, the relaxed mAP is higher than the strict mAP. Still, it is interesting to note that many models achieve a high relaxed mAP, while differences between models are more evident when considering the strict mAP. The top-ranked model in this track, SMatchToPR, is also ranked first when considering the average *value* of the mAP scores, as well as p@50%. However, the other models tend to swap positions when considering different measures, though there is a distinct group of high performing models. Figure 1(a) depicts strict mAP scores for each topic and stance, for the top performing models and for BarH. Evidently, some topic/stance pairs are easier than others for all models.

5.2 Generation Track

The results of the *Generation Track* are presented in Table 3 for both the matching aspect (mAP) as well as the quality of the generated KPs. The full results of the comparative annotation of the sets can be found in Table 6 in the appendix. In general, BarH clearly outperformed all submitted models, while SMatchToPR obtained the best results amongst the submitted models. Figure 1(b) shows the mAP scores for each topic and stance, for the top performing models and for BarH. We note that the inconsistent ranking of XLNet, namely first rank on relevance, while scoring much worse on other quality questions, stems from the fact that the model provided the set of KPs of the positive stance arguments for both the positive and negative stance. Although such a result can achieve high relevance, it fails when considering stance and redundancy, as indicated by the results.

It is interesting to note that BarH outperformed all submitted models by a significant margin, even though it had a relatively low rank in *Matching Track*. This could be due to its specific failure in matching to expert KPs in one of the topic/stance pairs (see Figure 1(a), second set of bars) coupled with the usage of macro-averaging.

When designing the task, we did not direct the participants regarding whether to use an extractive or abstractive approach for generating the KPs. Indeed, three models, BarH, SMatchToPR, and XLNet used an extractive approach, selecting specific arguments as KPs, while three models, Enigma, pa_vincent, and mozhiwen opted for an

abstractive approach, generating KPs de novo. Evidently, the extractive approach prevailed. Out of the three models taking the abstractive approach, only Enigma generated comprehensible KPs. The other two, received near zero scores in the matching evaluation, and hence, according to KPA-2021 terms, were not included in the KP set evaluation.

6 Additional Analysis

Accurate matching of arguments to KPs aims to properly reflect the prevalence of each KP in the data to facilitate more informed decision making. In this section we analyze how well different models were able to reproduce the presumably true distribution of KPs in the input data. To that end, we focus on the *Matching Track*, and consider the distribution of the expert KPs as the ground truth distribution. As can be seen in the top row of Figure 2(a), the five KPs provided by the expert for this topic and stance are not equally represented in the argument set. Nearly 40% of the arguments were matched by the annotators to "Social media regulation is required to deal with malicious users" (kp-1-9); around 30% were matched to "Social media regulation can help to deal with negative content" (kp-1-8); and so on. Correspondingly, for each model submitted in this track, we compare the distribution of KPs as predicted by its matching to the manually obtained ground truth distribution. To that end, we use the Jensen Shannon divergence (JSd). The smaller the divergence, the more accurate is the summary generated by a model (based on the 50% of the arguments that pass the threshold). We find that the average topic-stance JSd for the top 5 ranking models ranges from 0.152 to 0.238 (see Table 7 in the appendix for more details). Interestingly, the models which perform best in the *Matching Track* evaluation, i.e., have the highest mAP scores, are not those with the lowest JSd scores. This can also be seen in Figure 2(b). The rank of all models based on their JSd score is plotted vs. their rank in the matching task, for all topics and stances. Although it is evident that models which perform poorly on the matching task are also ranked low by JSd, the overall correlation is low. In other words, a model with sub-optimal matching capabilities may still perform relatively well in capturing the overall distribution of the KPs in the input data.

In the *Generation Track*, there is no single set of KPs, since each model generates its own set. It is

Model	Rank	Strict mAP (r)	Relaxed mAP (r)	Mean of mAP (r)	Strict p@50% (r)
SMatchToPR *	1	0.789 (1)	0.927 (4)	0.858 (1)	0.848 (1)
NLP@UIT *	2	0.746 (3)	0.93 (3)	0.838 (2)	0.827 (3)
ModrnTalk *	3	0.754 (2)	0.902 (6)	0.828 (4)	0.806 (5)
Enigma *	4	0.739 (5)	0.928 (4)	0.833 (3)	0.828 (2)
MatchTStm *	5	0.745 (4)	0.902 (6)	0.824 (5)	0.808 (4)
mozhiwen	5	0.681 (9)	0.948 (1)	0.814 (6)	0.784 (7)
fengdoudou	7	0.676 (10)	0.944 (2)	0.81 (7)	0.753 (13)
poledro *	8	0.717 (6)	0.901 (8)	0.809 (8)	0.799 (6)
Fibelkorn	9	0.706 (7)	0.872 (12)	0.789 (9)	0.779 (9)
XLNet *	10	0.684 (8)	0.869 (13)	0.777 (10)	0.765 (11)
RoBERTa-LC *	11	0.661 (13)	0.885 (10)	0.773 (11)	0.774 (10)
pa_vincent	12	0.596 (15)	0.893 (9)	0.744 (14)	0.689 (15)
BarH *	13	0.674 (11)	0.865 (14)	0.77 (12)	0.78 (8)
np-quadrat	13	0.601 (14)	0.877 (11)	0.739 (15)	0.729 (14)
niksss	15	0.672 (12)	0.857 (15)	0.764 (13)	0.759 (12)
SiamBERT *	16	0.555 (16)	0.729 (16)	0.642 (16)	0.678 (16)
yamenajjour	17	0.539 (17)	0.682 (17)	0.61 (17)	0.667 (17)
danielwaechteruni	18	0.39 (18)	0.601 (18)	0.495 (18)	0.511 (18)

Table 2: *Matching Track* results. Ranks on each measure are in brackets. In addition to the official evaluation metrics, two more scores are provided: the average value of the strict and relaxed mAP values, and p@50% for the strict view. Submissions that provided a descriptions are marked with (*).




Figure 1: AP per topic and stance for the top models and BarH. (a) *Matching Track*. (b) *Generation Track*. Missing bars mean a value of 0.

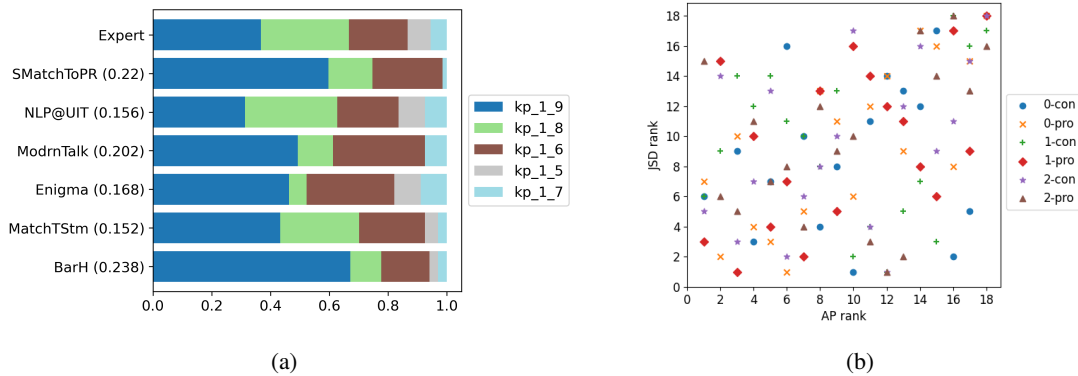


Figure 2: (a): The distribution of KPs among the arguments, in the pro stance on the topic *Social media platforms should be regulated by the government*, for the expert ground truth (strict view) and selected models, normalized to the number of arguments that match to a KP. The JSd vs. the expert distribution is given in brackets. (b) JSd rank vs. AP rank for all models in the *Matching Track*, per topic and stance. The map between topic id, and topic text can be found at Table 4. Rank 1 means best performance.

Model	Rank	mAP (r)	KP set rank
BarH	1	0.885 (1)	1 / (1,1,2)
SMatchToPR	2	0.818 (2)	2 / (2,1,2)
Enigma	3	0.491 (3)	4 / (2,4,3)
XLNet	4	0.443 (4)	3 / (4,3,1)
pa_vincent	5	0.096 (5)	None
mozhiwen	6	0 (6)	None

Table 3: *Generation Track* results. KP set rank shows the overall rank and in brackets the rank on each of the three questions: stance, redundancy, relevance.

Topic id	Text
0	Routine child vaccinations should be mandatory
1	Social media platforms should be regulated by the government
2	The USA is a good country to live in

Table 4: Topics of the test set.

therefore not relevant to compare the distributions over KPs as done above. However, it is possible to perform a qualitative anecdotal analysis. Table 8 in the appendix lists the KPs generated for one topic and stance, by the expert and the different models. Importantly, the expert produced KPs *a priori*, without seeing the actual set of arguments; while the models generated KPs based on the input arguments. Inspection of the generated KPs suggests that correspondingly, some KPs were perhaps overlooked by the expert. For example, BarH generated a KP concerning keeping schools safe, and another one concerning prevention of pandemic outbreaks, both of which are not in the expert set, even though there are expert KPs that are related. In Table 5 we show an example of one expert KP, to which we manually matched generated KPs from the top 3 models, as an anecdotal distribution comparison. In this example, BarH and SMatchToPR extracted the exact same KP, yet in BarH it maps to 35% of the arguments, which is in line with the weight of this KP in the expert ground truth, while SMatchToPR maps it to a much larger fraction of the arguments.

7 Discussion

We presented KPA-2021, the first shared task focused on key point analysis. The shared task received 22 submissions from 17 teams, covering both the *Matching Track* and the *Generation Track*. We presented the submitted models, their performance on different evaluation measures, and further analyzed the results.

As expected by its simpler nature, *Matching Track* received more submissions. However, ev-

idently, success in this intermediate matching task, does not guarantee success in the overall KPA task, which also requires generating the KPs. Future work should determine whether future KPA shared tasks should focus on the *Generation Track*, or perhaps modify the evaluation measures of the intermediate *Matching Track*, such that they better reflect the model performance in the full KPA task.

Given the inherent subjective nature of matching arguments to KPs, we opted for a ternary label, allowing argument-KP pairs to receive an "undecided" label, if the annotators votes were inconclusive. These undecided pairs are eventually considered either as positive or as negative matches, resulting with two gold standards, one potentially too strict, while the other perhaps too lenient. In future tasks, it may be valuable to consider a binary labeling scheme that will give rise to a single gold standard. Such labeling may be created conditionally, such that an undecided pair is marked as a positive match if and only if a minimum number of annotators marked it as positive and no other KP was labeled as matching to that argument.

A further point for future consideration, is the choice of a pre-defined threshold of 50%, that guided our evaluation. Although, this has the advantage of not requiring submissions to tune a confidence threshold, it has the limitation that it complicates the evaluation, since the ground truth coverage depends on the arguments, topic and stance. A possible alternative would be to require each model to set its own minimum confidence threshold to determine if an input argument should not be matched to any KP.

Finally, in this task, we used comparative evaluations to determine the quality of the sets of generated KPs. Obviously, this results with ranking-based evaluation. Establishing an absolute evaluation metric in this context merits further investigation.

We thank all participating teams for taking the time to participate in this challenging new shared task. We believe that these contributions, along with the data collected and shared in this report, will be valuable to further advance the research around KPA and related tasks.

Acknowledgments

We would like to thank Roy Bar-Haim for his helpful comments and suggestions.

Model	KP	Fraction
expert	Children should not suffer from preventable diseases	29.0
BarH	Vaccines help children grow up healthy and avoid dangerous diseases	35.0
SMatchToPR	Vaccines help children grow up healthy and avoid dangerous diseases	77.0
Enigma	Child vaccinations are necessary to protect against diseases in the child	50.0
	A mandatory vaccination is important to protect against diseases in the child	25.0

Table 5: KPs for *Routine child vaccinations should be mandatory* in the pro stance, related to the expert KP on the top row. Fraction is out of all arguments with a matched KP (for expert we use the strict view).

References

- Milad Alshomary, Timon Gurcke, Shahbaz Syed, Philipp Heinisch, Maximilian Spliethöver, Philipp Cimiano, Martin Pothast, and Henning Wachsmuth. 2021. Key point analysis via contrastive learning and extractive argument summarization. In *Proceedings of the 8th Workshop on Argument Mining*, Online. Association for Computational Linguistics.
- Reinald Kim Amplayo and Mirella Lapata. 2021. [Informative and controllable opinion summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2662–2672, Online. Association for Computational Linguistics.
- Roy Bar-Haim, Lilach Eden, Roni Friedman, Yoav Kantor, Dan Lahav, and Noam Slonim. 2020a. [From arguments to key points: Towards automatic argument summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4029–4039, Online. Association for Computational Linguistics.
- Roy Bar-Haim, Lilach Eden, Yoav Kantor, Roni Friedman, and Noam Slonim. 2021. [Every bite is an experience: Key Point Analysis of business reviews](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3376–3386, Online. Association for Computational Linguistics.
- Roy Bar-Haim, Yoav Kantor, Lilach Eden, Roni Friedman, Dan Lahav, and Noam Slonim. 2020b. [Quantitative argument summarization and beyond: Cross-domain key point analysis](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 39–49, Online. Association for Computational Linguistics.
- Tal Baumel, Raphael Cohen, and Michael Elhadad. 2016. Topic concentration in query focused summarization datasets. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, page 2573–2579. AAAI Press.
- Ralph Allan Bradley and Milton E. Terry. 1952. [RANK ANALYSIS OF INCOMPLETE BLOCK DESIGNS: THE METHOD OF PAIRED COMPARISONS](#). *Biometrika*, 39(3-4):324–345.
- Sihao Chen, Daniel Khashabi, Wenpeng Yin, Chris Callison-Burch, and Dan Roth. 2019. [Seeing things from a different angle: discovering diverse perspectives about claims](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 542–557, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emanuele Cosenza. 2021. Key point matching with transformers. In *Proceedings of the 8th Workshop on Argument Mining*, Online. Association for Computational Linguistics.
- H. Dang. 2005. Overview of duc 2005.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hady Elsahar, Maximin Coavoux, Jos Rozen, and Matthias Gallé. 2021. [Self-supervised and controlled multi-document opinion summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1646–1662, Online. Association for Computational Linguistics.
- Shai Gretz, Roni Friedman, Edo Cohen-Karlik, Asaf Toledo, Dan Lahav, Ranit Aharonov, and Noam Slonim. 2020. [A large-scale dataset for argument quality ranking: Construction and analysis](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7805–7813.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Manav Nitin Kapadnis, Sohan Patnaik, Siba Smarak Panigrahi, Varun Madhavan, and Abhilash Nandy. 2021. Team enigma at argmining-emnlp 2021: Leveraging pre-trained language models for key point matching. In *Proceedings of the 8th Workshop on Argument Mining*, Online. Association for Computational Linguistics.

- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [ALBERT: A lite BERT for self-supervised learning of language representations](#). *CoRR*, abs/1909.11942.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. [Generating wikipedia by summarizing long sequences](#). In *International Conference on Learning Representations*.
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Yao Lu, Yue Dong, and Laurent Charlin. 2020. [MultiXScience: A large-scale dataset for extreme multi-document summarization of scientific articles](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8068–8074, Online. Association for Computational Linguistics.
- Hoang Viet Phan, Long Tien Nguyen, Long Duc Nguyen, and Khanh Doan. 2021. Matching the statements: A simple and accurate model for key point analysis. In *Proceedings of the 8th Workshop on Argument Mining*, Online. Association for Computational Linguistics.
- Jan Heinrich Reimer, Thi Kim Hanh Luu, Max Henze, and Yamen Ajjour. 2021. Modern talking in key point analysis: Key point matching using pretrained encoders. In *Proceedings of the 8th Workshop on Argument Mining*, Online. Association for Computational Linguistics.
- Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer, and Ido Dagan. 2021. [Extending multi-document summarization evaluation to the interactive setting](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 657–677, Online. Association for Computational Linguistics.
- Shahbaz Syed, Roxanne El Baff, Johannes Kiesel, Khalid Al Khatib, Benno Stein, and Martin Potthast. 2020. [News editorials: Towards summarizing long argumentative texts](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5384–5396, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Assaf Toledo, Shai Gretz, Edo Cohen-Karlik, Roni Friedman, Elad Venezian, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. 2019. [Automatic argument quality assessment - new datasets and methods](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5625–5635, Hong Kong, China. Association for Computational Linguistics.
- Andrew Turpin and Falk Scholer. 2006. [User performance versus precision measures for simple search tasks](#). In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '06*, page 11–18, New York, NY, USA. Association for Computing Machinery.
- Lu Wang and Wang Ling. 2016. [Neural network-based abstract generation for opinions and arguments](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 47–57, San Diego, California. Association for Computational Linguistics.
- Yumo Xu and Mirella Lapata. 2020. [Coarse-to-fine query focused multi-document summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3632–3645, Online. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. [PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11328–11339. PMLR.

A Appendix

Model	Stance	Redundancy	Relevance
BarH	95.46	42.73	28.63
SMatchToPR	2.91	42.92	28.62
Enigma	1.37	4.47	0.0
XLNet	0.27	9.89	42.76

Table 6: Bradley Terry scores for the comparative evaluation of the top models in *Generation Track*.

Model	Routine child vaccinations should be mandatory		Social media platforms should be regulated by the government		The USA is a good country to live in		Average
	pos	neg	pos	neg	pos	neg	
SMatchToPR	0.29	0.13	0.247	0.245	0.157	0.252	0.22
NLP@UIT	0.328	0.144	0.0455	0.205	0.0585	0.153	0.156
ModrnTalk	0.238	0.131	0.242	0.216	0.125	0.263	0.202
Enigma	0.223	0.131	0.232	0.158	0.146	0.118	0.168
MatchTStm	0.181	0.105	0.0797	0.245	0.107	0.196	0.152
BarH	0.431	0.172	0.236	0.205	0.121	0.261	0.238

Table 7: Jansen-Shannon divergence for top five models and the state of the art model. The divergence is reported for each motion, and is also averaged. Best score for each model is bolded.

model	key point	%
Expert	Routine child vaccinations should be mandatory to prevent virus/disease spreading	37
	Children should not suffer from preventable diseases	29
	Child vaccination saves lives	20
	Routine child vaccinations are necessary to protect others	8
	Routine child vaccinations are effective	6
BarH	to keep schools safe children must be vaccinated	35
	Vaccines help children grow up healthy and avoid dangerous diseases	35
	The use of child vaccines saves lives	11
	our children have the right to be vaccinated	6
	in this way, pandemic outbreaks are avoided	5
	After getting vaccinated, our immune system produces antibodies.	4
SMatchToPR	Child vaccination should be mandatory to provide decent healthcare equally	2
	Vaccines help children grow up healthy and avoid dangerous diseases	77
	child vaccinations should be mandatory to provide decent health care to all.	12
Enigma	child vaccinations should be mandatory so our children will be safe and protected.	9
	Child vaccinations are necessary to protect against diseases in the child	50
	A mandatory vaccination is important to protect against diseases in the child	25
	Children should be vaccinated against diseases	23
XLNet	Children should be vaccinated against diseases as soon as possible	2
	It is important to protect the health of the next generation so to do that means the routine child vaccinations have to be mandatory.	57
	vaccines must be compulsory for children because in this way we prevent later diseases in infants	11
	Making routine child vaccinations mandatory would prevent the unnecessary deaths and suffering that currently result from the diseases that the vaccines prevent.	8
	child vaccinations should be mandatory so our children will be safe and protected.	8
	Child vaccination has to be mandatory since it takes care of our children to be immune to any type of disease so I think it has to be mandatory	4
	childhood vaccines should be mandatory because a child at an early age needs to strengthen their immune system and the state should be in the power to provide such protection	4
	If routine infallible vaccinations are mandatory, that way we avoid the spread and make our children immune from any disease.	2
The main reason is that many diseases have been prevented thanks to vaccines, thus millions of lives have been saved.	2	
pa_vincent	vaccines have be able to protect diseases	46
	child are be used to prevent diseases	33
	parents vaccines should be safe for children	7
	vaccines harms risk of the child	6
	vaccines is important for protection	6

Table 8: Expert and generated key points for *Routine child vaccinations should be mandatory* pro stance. Fraction is out of all arguments with a matched KP (for expert we use the strict view).