

slide 3: One of the first analysis we carry out is to check for any obvious trends for the accepted papers made by top affiliations to each conference. We consider a top affiliation one which had a large number of accepted papers at a conference in the last five years. The assumption is that for large conferences at least, the top 20 places each year will be taken by more prolific affiliations, likely to have participated in the past to the conference. It is unlikely that affiliations which have few sporadic research papers accepted in the last years are going to be present in the top 20 places. The plot shows the number of full research papers accepted at the KDD conference between 2011 and 2015 for the top 20 affiliations. The length of each line maps the range of the number of accepted papers for the affiliation and the mean number of papers is marked by the larger dot on each line. The plot shows the mean number of papers across all years could be a good predictor to how an affiliation will score in the future.

slide 4: In this first phase, we compute the probability of a full research paper is published by a given affiliation, based on the number of all their accepted full research papers between 2011-2015. The affiliations are ranked according to these probabilities and this represents our submission for the first phase of the competition. You can see in the table that predictions made using probabilities computed on more years (e.g. compute the probabilities over the entire range 2011-2014) are generally more accurate. They also follow an ascending trend for SIGIR when predicting the next year (e.g. predict 2015 rankings), than using 2011-2013 to compute the probabilities and predicting 2014 rankings. The predictions for 2013 using only the past couple of years are the worst across all the conferences.

slide 5: Our initial models use the number of accepted papers per year per affiliation as main predictor, but this predictor misses the fractional contributions of authors to the final affiliation rank. The relevance score, as calculated in the competition rules, distributes the paper score to each author and thus, each related affiliation receives only a fraction of the score. Starting with the second phase of the competition, we decide to use the relevance score as the prediction target in our models. So we train our models to predict the relevance directly.

slide 6: We check if the number of full research papers per conference per year is correlated with the total number of papers at the conference for that same year. Besides research papers, there are usually other types of contributions also present in the MAG such as journal papers, workshop papers or posters. We plot the number of papers on the Y axis across all the years on the X axis. The solid line represents the number of all the papers, while the dotted line is the number of full research papers only. The figure shows the two curves look similar for most the conferences and there is correlation between the number of full research papers and all the papers accepted at each conference. We aim to use two types of models, GBDT and mixed models. The first is more predictive, while the latter is more interpretable.

slide 7: We use this correlation to extend our dataset and consider all the papers for our predictions. We build the dataset as combinations of (conferences x affiliations x years). The large number of combinations helps grow the dataset from several hundreds samples to several thousands.

slide 8: We already know the mean historical relevance has a strong impact on how relevant an affiliation will be at a particular conference in the next year. Besides this, other features such as weighted trends, which are derived from past relevance values, as well as authors publishing trends also have an impact on the future relevance. First we compute simple statistics to measure different properties of past years' relevance scores. Then we try to capture the relevance trend across past years. Finally, another set of features is built using the author impact factor (AIF) measure. For a given author, the AIF is computed as the number of citations the author receives in a year for his publications from previous 5 years. Our final submission for the second phase are the prediction made by a single GBDT model using all the papers from 2011-2015 and the mentioned features.

slide 9: Most researchers publish their work at different conferences. However they specialize in a specific area and so the conferences they publish at have to be similar at least in a few respects. We use authors and keywords from the papers in MAG to cluster similar conferences together. It is a straightforward way to grow the dataset even more. The assumption is the information from related conferences will enforce the patterns discovered by the models, because prolific affiliations are prolific across all conferences they submit to, not just at one of them. We compute the Jaccard similarity for both authors and keywords for any pair of conferences in the MAG. From this, we can determine which conferences are for example most similar to KDD in terms of common authors and common papers' keywords. The table shows the most related conferences to KDD.

slide 10: As we shown previously, the number of full research papers is correlated with the total number of papers an affiliation has at a conference. Although the full research papers are not explicitly marked in the MAG before 2011, we assume this also applies to papers before this year. So we extend our dataset using papers from year 2000 and onward. The number of observations we use to train our models in this final phase reaches for some conferences almost 100K.

slide 11: In this phase, we keep most of the features we used in the previous competition phase, but we try to refine them. Besides the simple statistics which measure an affiliation's past relevance at a particular conference across all the years, we create four **window** versions of each of them for the last year up to four years back.

slide 12: We replace the features based on weighted trends from the previous phase with a set of more accurate time series based trend measures. We use the drift trend of historical relevance scores, which captures the increase or decrease of the relevance over time according to the average change in the past samples. This gives us an estimate of the relevance for the current year which we wrap in a new feature. Also we create five more features based on simple exponential smoothing and we experiment with different smoothing parameters. The new features are one-step-ahead forecasts, based on **all** past relevance scores. Depending on the values of the smoothing parameter, we give exponentially less weight to older observations, making newer observations more important and vice versa. Yet another new feature finds the smoothing parameter value to best fit the data points.

slide 13: Finally we tune our models configurations by combining all the features in feature sets together with different numbers of related conferences. We test our models against the baseline for each configuration. From all the configurations we choose for each of the conferences the one which outperforms the baseline across all years. You can see in the plot these results: the Y axis shows the NDCG@20 score obtained by the GBDT model (solid line) versus the baseline (dotted line) across all years marked on the X axis.