

Diverse Reviewer Suggestion for Extending Conference Program Committees

Christin Katharina Kreutz
Trier University
Germany
kreutzch@uni-trier.de

Krisztian Balog
University of Stavanger
Norway
krisztian.balog@uis.no

Ralf Schenkel
Trier University
Germany
schenkel@uni-trier.de

ABSTRACT

Automated reviewer recommendation for scientific conferences currently relies on the assumption that the program committee has the necessary expertise to handle all submissions. However, topical discrepancies between received submissions and reviewer candidates might lead to unreliable reviews or overburdening of reviewers, and may result in the rejection of high-quality papers. In this work, we present DiveRS, an explainable flow-based reviewer assignment approach, which automatically generates reviewer assignments as well as suggestions for extending the current program committee with new reviewer candidates. Our algorithm focuses on the diversity of the set of reviewers assigned to papers, which has been mostly disregarded in prior work. Specifically, we consider diversity in terms of professional background, location and seniority. Using two real world conference datasets for evaluation, we show that DiveRS improves diversity compared to both real assignments and a state-of-the-art flow-based reviewer assignment approach. Further, based on human assessments by former PC chairs, we find that DiveRS can effectively trade off some of the topical suitability in order to construct more diverse reviewer assignments.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; *Specialized information retrieval*.

KEYWORDS

Reviewer assignment, program committee extension, reviewer recommendation, reviewer coverage, flow-based algorithm

ACM Reference Format:

Christin Katharina Kreutz, Krisztian Balog, and Ralf Schenkel. 2021. Diverse Reviewer Suggestion for Extending Conference Program Committees. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI-IAT '21)*, December 14–17, 2021, ESSENDON, VIC, Australia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3486622.3493931>

1 INTRODUCTION

Scientific publishing heavily relies on peer review, which is typically performed by members of the program committee (PC) of a

conference. In general, PCs need to grow and change each year: to keep up with the increasing number of submissions [24], to avoid tunnel vision as well as unchanging perceptions of good or bad concepts [5] (e.g., the ACM SIGSOFT policy recommends to change one third of the members each year [24]), and former PC members might become unavailable [6]. According to current practice, organisers compose the PC before the submission period of manuscripts ends. Once submissions are closed, each manuscript gets a number of PC members, also called reviewers, assigned by the PC chairs, either manually or automatically (based on bidding information or preferred topics, entered by reviewers) [13, 21, 22]. Importantly, to the best of our knowledge, current approaches to reviewer recommendation assume a perfectly composed PC, and do not consider modification or extension as a necessity.

To the best of our knowledge, there is currently no way of reliably estimating the topical composition or amount of incoming submissions. Therefore, a previously disregarded problem is the possible mismatch between the expertise of current PC members and the expertise required for the assessment of all submissions. This problem may be further amplified by the ever-changing PC. Consequences of the mismatch might result in manuscripts tackling topics far from the PC's interests being less favourably reviewed [16] and a general overburdening of reviewers. This, in turn, might lead to innovative and complex submissions being rejected solely due to low-quality reviews [1] or failure to find errors in submissions [20].

A solution for the above issues would be the inclusion of new and additional PC members after the submission period ended, but before the review assignments have been made. This can especially help to cover new or emerging research topics [6] and to ensure that under-represented groups can gain exposure and reviewing experience [19]. Identification of appropriate candidates is challenging as PC members should be diverse in localities, seniority [16, 20], research topics and gender [16]. Furthermore, suggested candidates should be explainable, in order to aid the conference chairs in effective and efficient decision making.

In this paper we focus not only on the automatic assignments of reviewers to submissions (i.e., *reviewer assignment*), but also introduce and address the problem of *reviewer coverage*: ensuring the assignment of suitable reviewers to all submissions. This gives rise to the novel task of *reviewer suggestion* for PC extension: given the current PC and all submitted manuscripts of a venue, recommend new reviewer candidates to be added to the PC. Note that these two tasks are interconnected: our reviewer assignment method identifies submissions that would not receive adequate reviewers using the current PC, which in turn triggers the suggestion of new reviewers to extend the PC. Those newly included persons should not only be capable of ideally assessing multiple manuscripts but

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WI-IAT '21, December 14–17, 2021, ESSENDON, VIC, Australia

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9115-3/21/12...\$15.00

<https://doi.org/10.1145/3486622.3493931>

also ensure diversity of the whole PC. Note that some gaps in the PC can be identified without requiring paper-reviewer assignments (e.g., not enough senior reviewers, reviewers from a given location or stark imbalance in academic vs. non-academic backgrounds of reviewers), while other gaps may only be identified once a (preliminary) assignment is done.

The main contribution of this work is a flow-based reviewer suggestion and PC extension approach, termed *DiveRS*. The main idea behind *DiveRS* is to iteratively identify submissions that are unlikely to get a set of suitable reviewers assigned. These problematic submissions and currently underrepresented diversity aspects (professional background, location or seniority) determine the reviewer candidates for inclusion in the PC to support a feasible reviewer assignment. We capture these characteristics in a constrained optimisation problem. At its core, *DiveRS* relies on a reviewer assignment method, which considers reviewers as a set for each paper, in order to satisfy diversity constraints. Additionally, reviewers' individual upper and bounds of the numbers of papers to review, and their conflicts of interests, also need to be respected.

We evaluate *DiveRS* on real-world conference datasets in two parts. First, we compare it on the task of *reviewer assignment* against the current state-of-the-art, PR4All [21], and against real assignments, in terms of both established measures (mean number of papers assigned, fairness, and textual diversity of reviewer sets) as well as novel measures (diversity and dependency between reviewers). We show that *DiveRS* achieves fairness that is on par with PR4All, while being superior in terms of diversity. Second, we evaluate the *reviewer suggestion* task by asking actual PC chairs to assess the generated suggestions for PC extension in terms of relevance, usefulness, and accompanying explanation. Our results indicate that *DiveRS* can effectively trade off topical suitability in order to improve the diversity of the assigned reviewer sets.

In summary, we make the following contributions:

- We propose the reviewer coverage problem as an extension of the reviewer assignment problem, where we no longer assume the current PC to be perfectly suitable for all submissions. We define the extension of the PC, to accommodate possibly ill-covered submissions, as part of the objective.
- We present *DiveRS*¹, a novel reviewer assignment and PC extension approach. It incorporates previously overlooked diversity aspects in terms of professional background, location and seniority of reviewer candidates directly in the assignment process, and generates explainable suggestions for extending the PC.
- We propose new measures for evaluating the diversity and dependency of reviewer sets.
- We automatically evaluate our approach on two real-world datasets and demonstrate its suitability in manual evaluations with the actual PC chairs of these conferences.

2 RELATED WORK

Areas related to our work are *reviewer assignment*, which corresponds to the typical reviewer assignment problem, as well as the general field of *program committee construction*, which relates to the extension of PCs. For conference organisers there are many systems supporting the bidding and reviewer assignment process

but “[e]xtending PCs based on submitted papers” as identified as a future objective by Price and Flach [17] has not yet been tackled to the best of our knowledge. There have been efforts to expand expert sets to hold more persons similar to the ones already contained in the set [25] but these approaches differ from our research objective: instead of finding more similar experts, our goal is to suggest an unbiased and diverse set of reviewer candidates to better cover the topical composition of incoming submissions.

Reviewer Assignment. There is a multiplicity of author-topic models to capture topical relationships between authors and (their) papers [8, 9, 15, 18, 23]. We refrain from discussing them in detail or utilising them here, as our focus within assigning reviewers to submissions lies not only on topical similarity of the two, but more on diversity aspects.

Conry et al. [4] tackle the reviewer assignment problem with given bidding information as an optimisation problem with global criteria. They extend bidding data by predicting new preferences of reviewers, and utilise manuscript as well as reviewer similarities. Liu et al. [13] recommend n reviewers for each manuscript which are dependent on each other. They model reviewers' expertise, authority and diversity in a graph, which they traverse with random walk with restart. The number of co-authorships is modelled as authority. Tang et al. [22] propose a constraint-based optimisation framework that proposes sets of reviewers for query manuscripts and user feedback, if available. They incorporate expertise matching, authority aspects based on seniority, load balance and aim to maximise the topic coverage between reviewer sets and manuscripts, using LDA. Long et al. [14] study topic coverage and fairness of manuscript-reviewer assignments. They maximise the numbers of different topics of manuscripts in which the assigned reviewer set is knowledgeable. Additionally, they define and regard the influence of different conflict of interest types, such as the competitor relationship, in the assignment. Kou et al. [11] build upon [14] and instead observe a weighted topic coverage score. Their approach calculates the assignment resulting in the approximate maximum weight-coverage group-based scores, while fulfilling workload and reviewer set size constraints.

Jecmen et al. [7] provide a solution for the reviewer assignment problem, which focuses on supporting the integrity of the peer review process. The approach prevents reviewers' manipulation efforts in the assignment to either submit overly positive or negative feedback as well as de-anonymise the reviewing process. Here, the similarity between manuscripts and reviewers' profiles (expertise) is a critical factor in the randomised assignment. Kobren et al. [10] introduce a paper-reviewer-assignment strategy which incorporates upper and lower load bounds per reviewer, guarantees a minimal required expertise in the area of the submission from all assigned reviewers and optimises a global objective. They present a linear programming and min-cost flow-based heuristic approach.

The Toronto Paper Matching System (TPMS) [3] conducts automatic reviewer assignment for all manuscripts submitted to a conference by using either word count representation or LDA topics, but can also incorporate reviewers' bids on submissions. TPMS supports some constraints: papers must be reviewed by three reviewers, and reviewers are assigned not more than a certain limit of papers. Reviewers for manuscripts are determined based on expertise extracted from their published papers and maximising the

¹*DiveRS* implementation: <https://github.com/kreutzch/DiveRS>

similarity between reviewers and manuscripts. Stelmakh et al. [21] use TPMS in PR4All; they propose an approach utilising a max-flow algorithm to identify the top- k papers submitted to conferences, which should be accepted. They focus on fairly assigning suitable reviewer sets to all submissions via TPMS, especially those which received low similarity with all reviewer candidates. This approach is considered as the state of the art for flow-based reviewer assignment [10].

We note that the datasets used in related work are mostly not available online and even fewer contain all submissions of a conference, i.e., include rejected papers. Those that remain either do not contain the real reviewers (ICLR 2018 [7, 21]) or do not contain names of both reviewers and authors (MIDL, CVPR and CVPR2018 [10]). Thus, to the best of our knowledge, there is no publicly available dataset including rejected papers, and non-anonymised reviewer and author names from a real conference. Therefore, we create our own datasets based on real conference data in §5.1.

Program Committee Construction. Han et al. [6] recommend PC members for conferences based on the previous year’s PC and core authors, preferring candidates socially close to current chairs. They build a language model for a conference by aggregating previously published papers and compare it to PC candidates’ publications. Authoritativeness of candidates influences the recommendations. Sekar [19] introduces EZ-PC, a tool to define constraining factors and help automate the PC formation process as an integer linear programming problem. Several factors are considered: topical coverage, diversity of the PC, avoiding over-representation of groups and keeping the PC size manageable. The main differences between their work and ours are that diversity constraints in EZ-PC are on the PC level, and they do not support reviewer assignment.

3 PROBLEM SETTING

3.1 Problem Statement

We define the *reviewer coverage problem* (RCP) as an extension of the *reviewer assignment problem* (RAP) for scientific manuscripts. Both problems have the underlying goal of finding suitable *sets of reviewers* for each manuscript. These sets need to be constructed such that (i) reviewer expertise is sufficient for the topics of the respective manuscript, (ii) there are no conflicts of interests between authors of submissions and reviewers, and (iii) overall reviewer load constraints are met. Contrasting with RAP, RCP does not assume that the current PC is perfect (i.e., has sufficient coverage), but explicitly allows for its extension by adding reviewer candidates from an *extended reviewer candidate pool* (ERC). So the immediate goal for RCP is the suggestion of new PC members, which leads to sufficient reviewer expertise for all submissions, while also ensuring diversity in the PC in terms of (i) seniority, (ii) location, and (iii) industrial/academic affiliation.² An additional desirable condition for the inclusion of new PC members is their ability to review multiple papers. Formally, the output of RCP is twofold: (1) a ranked list of reviewer suggestions to include in the PC and (2) an assignment of reviewer sets to submissions.

²Gender would also be a desirable diversity aspect for PCs [16], but we consciously refrain from touching this subject due to the challenges involved in collecting potentially personal information from reviewers for inclusion in our datasets.

3.2 Notation

M describes the set of submissions to a conference for which reviewers from the program committee PC need to be assigned. A single reviewer is addressed as $r_i, i \in \{0, \dots, |PC| - 1\}$ or only by their index i . We address a single submission as $m_j, j \in \{0, \dots, |M| - 1\}$ or only by their index j . An assignment is *feasible* if all submissions are assigned a predefined number of reviewers λ , the number of submissions a reviewer is assigned lies between a predefined lower (μ_i^l) and upper bound (μ_i^u), which is specific for each reviewer i , and conflicts of interests (COI) are not violated by the assignment. The reviewer set assigned to a submission j under a feasible assignment A is denoted by $R_A(j)$. We store similarities of reviewers and submissions in $S \in [0, 1]^{|PC| \times |M|}$; the similarity S_{ij} of reviewer i with submission j is seen as a proxy for expected review quality [21] and can be determined, e.g., by the cosine similarity between TF-IDF representations of j ’s and i ’s profiles, composed of their papers. In case of a COI between i and j , we set $S_{ij} = -1$. We store dependencies between reviewers in $dep \in \{0, 1\}^{|PC| \times |PC|}$; dependencies such as recent co-authorships between reviewers i and k are expressed by $dep_{ik} = 1$ if there is a dependency and 0 otherwise.

4 METHOD

We introduce **DiveRS**, a **Diverse Reviewer Suggestion** system for extending conference program committees. It focuses not only on fairness of reviewer assignments but also considers diversity in professional background, location of reviewer candidates and their seniority. We build on and extend a previous state-of-the-art flow-based approach [21], by explicitly modelling diversity as a layer in the flow-graph; see Fig. 1.

4.1 Modelling Diversity

We focus on diversity in three different areas: professional background, location and seniority. We integrate these properties of the assignment in a specific layer in our flow network between papers and reviewers (diversity layer **L4** in Fig. 1). Diversity in *professional background* means that each reviewer set has to contain at least one reviewer working in academia and one reviewer (possibly the same one) working in industry. For diversity in *location* it would be desirable to include reviewers in a reviewer set with locations from completely different geographical locations. The goal here is to not have all reviewers in a set being located on the same continent. We achieve diversity in *seniority* by enforcing each reviewer set to contain at least one senior researcher [3, 22]. Meanwhile, overburdening of reviewers from underrepresented backgrounds can be prevented by decreasing their possible reviewing load. Satisfying all diversity constraints might lead to an increase of the PC size.

4.2 Algorithm

DiveRS identifies submissions with high probability of not obtaining enough suitable (topically fitting and diverse from each other) reviewers and adds new reviewers to the PC accordingly. It then constructs suitable reviewer sets for all submissions from the extended PC. Our reviewer suggestion approach is inspired by PR4All [21], the current state-of-the-art in flow-based reviewer assignment [10]. However, PR4All tackles the reviewer assignment problem (RAP), which is only one element of the larger reviewer coverage problem

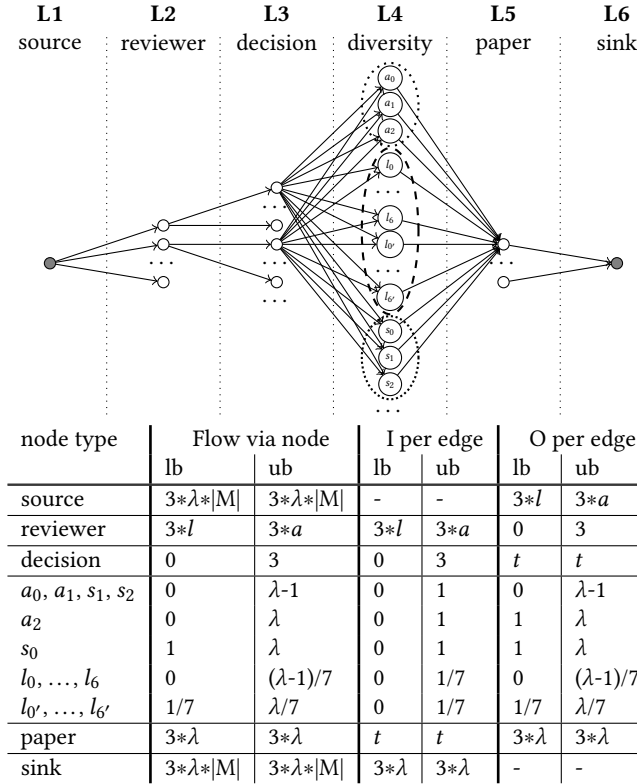


Figure 1: Top: A simplified version of the flow network constructed by DiveRS. Only the depicted edges between neighbouring layers allow flow. Background nodes in the dotted ellipse are used to ensure diversity in the professional background, those in the dashed ellipse are used to enforce diversity in the continent of the assigned reviewers and those in the densely dotted ellipse guarantee diversity in seniority. Bottom: Lower (lb) and upper bounds (ub) of incoming (I) and outgoing flow (O) per edge as well as the general flow via a specific node type with ability a , demand λ , lowest load l and amount of flow depending on the node type t .

(RCP) that we are addressing (cf. §3.1). We do not only construct suitable assignments but also identify possibly problematic papers and actively extend the PC to ensure diverse reviewer sets.

We first discuss the limitations of PR4All in §4.2.1, followed by the introduction of DiveRS’s reviewer assignment subroutine in §4.2.2 and its main routine in §4.2.3 which is responsible for identifying problematic submissions and suitable reviewer candidates.

4.2.1 PR4All. The goal of PR4All [21] is the fair assignment of suitable reviewer sets for all submissions with a focus on the most disadvantaged ones. The iterative approach fixes one reviewer set for the worst off submission in each iteration. Each iteration constructs partial reviewer sets for all unassigned submissions consisting of the most similar reviewers. This is their central optimisation problem. The partial sets are merged and considered a possible assignment. One assignment resulting in the highest fairness is computed out of several of these possible assignments. From the

best overall merged assignment, the worst off paper is finally assigned its reviewers. Fixed (worst off) papers are disregarded in the next iterative assignment and merge steps until all papers are fixed.

Due to the merge step, PR4All cannot introduce new conditions for the single reviewers and reviewer sets on the final level only, e.g., lower bounds (μ^l) for the number of assigned submissions for each reviewer or that each set must contain at least one reviewer from industry and one from academia. Instead, these lower bounds for reviewers and conditions for reviewer sets would be applied during all parts of the assignment process. Overcoming this issue is non-trivial as all partial assignments which are then merged fulfilling the new conditions could also lead to violated upper bounds (μ^u) and an excess of industry reviewers per final reviewer set. For their initial run with sets of size 1, the one reviewer would be required to represent both professional backgrounds which is hard to find. Conditions that only merged assignments have to fulfil cannot be realised in the described optimisation problem. So, PR4All prevents definition of desirable properties for final assignments that surpass mere similarity, such as diversity in certain properties.

4.2.2 DiveRS Subroutine: Reviewer Assignment. We strive to overcome some of the weaknesses in reviewer assignment encountered in PR4All: we introduce individual upper (μ_i^u) and lower bounds (μ_i^l) of reviewing abilities for each reviewer i [10]. The lower bound describes the number of submissions, a reviewer has to review at least. Additionally, we allow for the definition of dependencies between reviewers (e.g., in case of shared current affiliations or recent collaborations), as reviewers in a set should have distinct affiliations to make sure that their opinions are sufficiently independent from each other [7]. The resulting constraint can mathematically be described by the expression $con_I : \sum_{j \in M} \sum_{i, k \in R_A(j), i \neq k} dep_{ik} = 0$.

Our goal is to assign reviewers to the best fitting submissions to maximise the overall similarity between assigned reviewers and submissions. The following equation formulates the optimisation objective: $maximise J_f^S(A) := \sum_{i \in R_A(j), j \in M} f(S_{ij})$ while all submissions receive λ reviewers, dependencies between reviewers, COIs, diversity constraints of reviewer sets as well as reviewers’ lower and upper abilities are not violated. f is a monotonically increasing function used to transform similarity values $[0 : 1] \rightarrow [0 : \infty]$ [21].

Algorithm 1 (main routine) and Algorithm 2 (subroutine) depict the pseudo code of our approach. In the subroutine we construct our flow network such that reviewers review a number of submissions limited by their upper and lower bounds. Submissions are reviewed by λ reviewers. Each reviewer set for a submission is diverse in professional background (at least one from industry and one from academia), location (not all from the same continent) and seniority (at least one senior reviewer). We only allow the allocation of reviewers to submissions, if this combination is contained in *pairs*. The decision if a reviewer is assigned to a submission is contained in L3, if there is flow over an edge $(i, i * |M| + j)$ between L2 (reviewer i) and L3 (decision to review submission j), i is assigned as reviewer for j . If we can compute a max flow, we find a feasible assignment.

4.2.3 DiveRS Main Routine: Reviewer Suggestion for PC Extension. In the main routine we generally first check if the original PC contains enough reviewers such that each submission can be assigned

Algorithm 1 DiveRS main routine: reviewer suggestion for PC extension.

Input: $\lambda, M, PC, S, \mu^l, \mu^u, dep, acaInd, location, seniority, tries, S_ERC, \mu^l_ERC, \mu^u_ERC, dep_ERC, acaInd_ERC, location_ERC, seniority_ERC, \theta, \kappa$

Output: Reviewer assignment A , problematic papers $M_{<\theta}$

- 1: while PC is not able to produce assignment based on ability and seniority or professional background and $|ERC| > 0$: include new reviewers from underrepresented aspects with highest average similarities to all manuscripts
 - 2: if abilities of PC are not enough to find assignment: terminate with error
 - 3: \forall reviewer-submission pairs from S and S_{ERC} set similarity = -1 if similarity $< \theta$ (equivalent to COI)
 - 4: $M_\theta = M$ w/o submissions with all similarities $< \theta$
 - 5: delete reviewers r from PC where $\mu_r^l >$ number of submissions with which they have similarity $\geq \theta$
 - 6: $pairs =$ compute all pairs of reviewers in PC and papers in M_θ
 - 7: **while** $sub(\lambda, M_\theta, PC, S, [0]^{|PC|}, \mu^u, dep, acaInd, location, seniority, pairs)$ does not produce assignment **do**
 - 8: $fewCandidatePapers =$ papers with $< \lambda$ reviewers w/o COI
 - 9: run sub multiple times w/o $fewCandidatePapers$ and w/o predefined % of submissions to identify (possibly problematic) submissions where run fails, i.e., for which no assignment can be computed due to ill-fitting or few reviewers in the submission's area; adjust $pairs$ and M_θ for runs, papers with highest probability of failed run are $problemPapers$
 - 10: insert up to κ reviewers in PC from ERC fitting $fewCandidatePapers +$ most problematic $problemPapers$ and underrepresented background variables best
 - 11: delete papers from M as out of scope for which not enough reviewer ($< \lambda$) candidates with similarity $\geq \theta$ can be found
 - 12: $pairs =$ compute all pairs of reviewers and papers
 - 13: **end while**
 - 14: $f_A = []$ // list of all feasible assignments
 - 15: **for** $try = 0, try \leq tries, try ++$ **do**
 - 16: $pairs_c =$ drop predefined percentage of $pairs$
 - 17: $A_{curr} = sub(\lambda, M_\theta, PC, S, \mu^l, \mu^u, dep, acaInd, location, seniority, pairs_c)$
 - 18: if $A_{curr} \neq \emptyset : f_A.append(A_{curr})$
 - 19: **end for**
 - 20: return most diverse assignment from $f_A, M - M_\theta$
-

someone from both professional backgrounds as well as one senior reviewer. Otherwise, we include new reviewers with missing diversity properties in the PC from ERC (l. 1). The ERC could, e.g., be composed of authors of former instances of a conference. The similarity threshold θ heavily influences DiveRS, it defines the minimal similarity between submissions and assigned reviewers [10] (l. 3-5, 11). If $\theta=0$, the algorithm often finds a solution to the reviewer assignment problem, computed by the subroutine after including new reviewers based on underrepresented diversity aspects (l. 7), and does not need to identify possibly problematic papers (l. 8-9). Those problematic papers (l. 9) are submissions which have a high probability of not getting assigned reviewers (i.e., where runs of the sub-routine oftentimes fail if they are part of M_θ). The higher

Algorithm 2 DiveRS subroutine: reviewer assignment step *sub*.

Input: $\lambda, M, PC, S, \mu^l, \mu^u, dep, acaInd, location, seniority, pairs, \theta$

Output: Computed reviewer assignment A_{curr}, \emptyset if unfeasible

- 1: **Initialization:** flow network (see Figure 1):
 - L1 (source, 1 vertex)
 - L2 (reviewers, vertex $\forall i \in PC$)
 - L3 (reviewer paper decision, vertex $\forall j \in M * \forall i \in PC$)
 - L4 (diversity, 3 vertex types, vertex $\forall x \in M * 20$, see Figure 1)
 - L5 (papers, vertex $\forall j \in M$)
 - L6 (sink, 1 vertex)
 - 2: Reset flow constraints for all vertices in the network: source, reviewer, decision, diversity, papers, sink
 - 3: $\forall (ij) \in pairs$: insert edge (i, j) between L2 and L3 (i.e., set $capacity[i, i * |M| + j] = 3$), adjust flow constraints
 - 4: Compute max flow, create assignment A_{curr} corresponding to max flow $\forall (ij)$: **if** flow on edge $(i, i * |M| + j)$ between L2 and L3 **then** assign reviewer i to submission j
 - 5: return A_{curr}
-

the value of θ , the more difficult it is to find a feasible assignment. A value κ describes the number of fitting inserted reviewers per iteration (l. 10). If a feasible assignment (l. 7) has been found for a reviewer set, we randomly exclude reviewers from reviewing submissions in order to find the most diverse assignment (l. 15-20).

Figure 1 depicts our network and associated flow constraints. Nodes a_x indicate the professional background of a reviewer ($a_0 =$ industry, $a_1 =$ academia, $a_2 =$ both). The different l_x indicate the location background of reviewers (l_y indicates the presence of a continent in continents associated with a specific reviewer while l_y' indicates the continent's absence in their continents; $l_0 =$ South America, $l_1 =$ Africa, $l_2 =$ Antarctica, $l_3 =$ Asia, $l_4 =$ Oceania, $l_5 =$ North America, $l_6 =$ Europe). Nodes s_x indicate the different levels of seniority of researchers ($s_0 =$ senior, $s_1 =$ advanced, $s_2 =$ junior).

In our implementation, we utilise Gurobi³, a commonly used [10, 21] solver software for mathematical optimisation.

4.3 Practical Issues and Effects of Parameters

Our approach tackles several practical issues which arise in PC extension and reviewer assignment:

- Submissions for which no suitable reviewers can be found as their topics might be out of scope of a current conference can be identified and considered manually.
- Reviewers that are part of the original PC should all be assigned at least one submission out of courtesy, even if they might no longer fit the topical composition of the conference. For subsequent instances of the conference, these individuals may no longer be invited to the PC. Our approach identifies such reviewers and is able to assign them to current submissions nevertheless.

Running time constraints influence the choice of parameters:

- The higher the similarity threshold θ is set, the more iterations (l. 7-12 A. 1) are required until a feasible assignment is found. The higher the number of included new reviewers per run κ is set (l. 10 A. 1), the longer one single run of the assignment step (A. 2) takes but in total less iterations might be needed. If κ is

³<https://www.gurobi.com/>

high, the total review load will be distributed among the many new candidates. In order to keep the PC comparably small, we advise to have a low κ and more iterations in total.

- The higher the bias towards incorporation of reviewers with underrepresented background variables (l. 10 A. 1), the less focus is put on similarity of reviewers and submissions. In consequence, fairness of assignments decreases while diversity increases.

Note that we do not separately handle sub- or meta-reviewing but DiverS can be used in these steps with different parametrisation.

5 EXPERIMENTAL SETUP

We present our experimental setup, introducing two new datasets (§5.1), our parameter settings (§5.2), an overview of established measures (§5.3), and novel ones for reviewer assignment assessment, namely diversity and dependency (§5.4).

5.1 Datasets

We evaluate on two real-world conference datasets based on the International Conference on the Theory of Information Retrieval (ICTIR) in 2019 (I'19) and 2020 (I'20). The data was made available to us by the conference organisers upon request and signing an NDA. The datasets include all manuscripts submitted to the conferences, not only the accepted ones, authors of submissions, reviewers, real reviewer-submission assignments, and constructed extended reviewer candidate pools. I'19 (I'20) contains 78 (65) papers submitted by 201 (184) authors. 43 (30) papers were accepted and 36 (35) rejected. There were 43 (67) reviewers. The extended reviewer candidate pool consists of 6,445 (5,692) authors from papers which appeared in CIKM, ECIR, ICTIR and SIGIR in the previous five instances of the conferences.

For all *reviewers*, we retrieved their DBLP key [12], COIs and dependencies (collaborators from the previous five years and persons with current shared affiliations), seniority, location, current affiliation(s) as well as information on whether they are working in industry and/or academia. Demographics were automatically derived from their affiliations and earliest published paper. Additionally, we collected the titles of their publications up until the year of the conference, and abstracts from the previous five years for papers which appeared with Springer or ACM. We performed further manual post-processing to ensure high data quality.

5.2 Parameter Settings

We assign each submission to three reviewers, following the practice of the I'19 and I'20 conferences. Similarity between submissions and reviewers (a concatenation of their publications' titles and abstracts) is taken to be the cosine similarity of TF-IDF-weighted document representations, thus all similarity values lie in $[0,1]$. We utilise $f(S_{ij}) = \frac{1}{1-S_{ij}}$ if $S_{ij} < 1$ and $1 * e^6$ otherwise [21]. For DiverS we set $tries = 25$, $\kappa = 10$ and $\theta > 0$ to .25 for I'20 and .15 for I'19.⁴ We set $\mu^u = 9$ for I'19 and $= 7$ for I'20 according to the real number of maximal assigned submissions per reviewer candidate.

For the manual evaluations we obtain reliable human assessments by asking respective PC chairs (3 from I'19 and 2 from I'20) to fill out a questionnaire.

⁴Different values for θ had to be chosen to find feasible assignments, as θ is highly dependent on topical fit between the submissions and the PC.

5.3 Established Measures

The following established measures describe the *quality* of reviewer assignments: mean *number of papers* assigned to single reviewers [10], *fairness* of the assignment [10, 21], and *average textual diversity* of reviewer sets [13].

Fairness of an assignment A is defined as the minimal summed similarity between any submission j and its reviewers $R_A(j)$ [21]: $\Gamma_f^S(A) = \min_{j \in M} \left(\sum_{i \in R_A(j)} f(S_{ij}) \right)$, with f being a monotonically increasing function $[0, 1] \rightarrow [0, \infty]$. *Average textual diversity* of reviewer sets is calculated by the average Kullback-Leibler (KL) divergence between pairs of reviewers assigned to the submissions [13]: $KL(A) = \text{avg}_{j \in M} \left(\frac{\sum_{i,k \in R_A(j), i \neq k} KL_Divergence(i,k)}{|R_A(j)| * (|R_A(j)| - 1) / 2} \right)$. We calculate this value on an unigram language model of the reviewer's publication information. Higher values for average KL-divergence indicate less similar reviewers in reviewer sets. Desirable complementary reviewers [3] produce a high value.

5.4 Novel Measures

We present a novel measure for quantifying the *diversity* of backgrounds of reviewers. We define diversity for reviewers that are part of a feasible assignment A , as a linear combination of background-, location-, and seniority-based diversity scores (each in $[0, 1]$): $Div(A) = \text{avg}_{j \in M} (Div_{BG}(j) + Div_L(j) + Div_S(j))$. Diversity can take values in $[0, 3]$, where higher values are more desirable. Note that diversity of one single reviewer set $R_A(j)$ can be computed using the same formula by setting $M = \{j\}$.

The component-level diversity scores are estimated as:

$$Div_{BG}(j) = 1 - \frac{|\sum_{i \in R_A(j)} profBG[i]|}{\lambda}$$

$$Div_L(j) = 1 - \frac{1}{\binom{2}{\lambda}} * \sum_{i,k \in R_A(j), i \neq k} \frac{|location[i] \cap location[k]|}{|location[i] \cup location[k]|}$$

$$Div_S(j) = \sum_{val \in \{0,1,2\}} \mathbb{1}(\exists i \in R_A(j) : seniority[i] == val) * \frac{1}{3},$$

where for each reviewer i , $profBG[i]$ indicates the professional background (0 if both, -1 if industry, 1 if academia), $location[i]$ denotes the distinct locations associated with i , and $seniority[i]$ describing the seniority level (0 if senior, 1 if advanced, 2 if junior).

We further quantify the *dependency* of an assignment as the percentage of reviewer sets with violated dependencies between reviewers i, k : $Dep(A) = \frac{\sum_{j \in M} \mathbb{1}(\exists i,k \in R_A(j) : i \neq k, dep_{ik} = 1)}{|M|} * 100$.

Example. Given: $M = \{j\}$, $R_A(j) = \{i \text{ (both, senior), } k \text{ (academia, senior)}\}$, i and k from different locations, $dep_{ik} = 0$. We can then compute $Div(A) = (1 - \frac{1}{2}) + (1 - \frac{1}{1} * 0) + (\frac{1}{3}) = \frac{11}{6}$ and $Dep(A) = 0$.

6 EXPERIMENTS

Recall that the output of RCP is twofold: (1) an assignment of reviewer sets to submissions and (2) a ranked list of reviewer suggestions to include in the PC. We thus divide our evaluation into two parts: an examination of reviewer assignments in §6.1, using both automatic (§6.1.1) and manual evaluation (§6.1.2), followed by an evaluation of reviewer suggestions using human assessments by the respective PC chairs in §6.2.

Table 1: Reviewer assignment results for the automatic evaluation in terms of mean workload per reviewer (mW/R) and all initial PC members (/PC), number of unused initial PC members (U) as well as dependency (Dep), fairness (Γ_f^S), average textual diversity (KL), and diversity (Div) of assignments per dataset and method. Methods marked with * correspond to the restrictive setting.

method	d.set	mW/R (/PC)	U	Dep	Γ_f^S	KL	Div
real	Γ 19	6.16 (5.44)	5	15.38	2.12	.45	1.51
PR4All	Γ 19	7.09 (5.44)	10	48.72	3.51	.52	1.58
$D_{\theta=0}$	Γ 19	6.69 (5.09)	11	0	3.31	.46	2.16
$D_{\theta=0}^*$	Γ 19	5.09 (5.09)	0	0	3.07	.45	2.13
$D_{\theta>0}$	Γ 19	6.16 (4.78)	11	0	3.68	.45	2.15
$D_{\theta>0}^*$	Γ 19	4.98 (4.78)	2	0	3.68	.45	2.13
real	Γ 20	3.73 (3.12)	11	24.62	2.4	.45	1.57
PR4All	Γ 20	4.88 (2.91)	27	47.69	3.62	.52	1.55
$D_{\theta=0}$	Γ 20	4.53 (2.87)	25	0	3.5	.47	2.04
$D_{\theta=0}^*$	Γ 20	2.87 (2.87)	0	0	3.18	.47	2.05
$D_{\theta>0}$	Γ 20	3.16 (2.03)	32	0	4.05	.44	2.12
$D_{\theta>0}^*$	Γ 20	2.23 (2.13)	4	0	4.05	.44	2.09

6.1 Part 1: Reviewer Assignment

For evaluating the reviewer set construction properties of our approach (conducted by our subroutine in §4.2.2) we compare different variants of our DiveRS (D_θ) algorithm against (1) assignments produced by a state-of-the-art flow-based reviewer assignment system, PR4ALL [21], and (2) the real reviewer assignments.

6.1.1 Automatic Evaluation. In our automatic evaluation, we report the established measures for reviewer assignment from §5.3, the newly introduced measures from §5.4, and the number of unused reviewers from the original PC.

In addition to the DiveRS default setting, we also report on a *restrictive setting*, where each reviewer i from the original PC who can review at least one submission (i.e., similarity $\geq \theta$) needs to be used in the final assignment ($\mu_i^l = 1$). This setting is desirable to prevent displeasing reviewers who have already been invited to the PC by not assigning them to a submission. In PR4All such an option is not given, including a lower bound for numbers of assignments is impossible as the approach merges assignment sets.

Table 1 reports the results of the automatic evaluation. DiveRS achieves the highest diversity scores regardless of the setting. Real assignments are worse in fairness and diversity than the automatically constructed sets. Usage of $D_{\theta>0}$ leads to fairer and mostly more diverse results compared to the $D_{\theta=0}$ variants. KL-divergence does not seem to change much between configurations, but PR4All produces sets with the highest score. Introduction of new PC members naturally reduces the mean workload per reviewer. With the restrictive DiveRS variants to include all reviewers from the original PC in the assignment (marked with *), we achieve fairness, KL, and diversity values comparable to the unrestricted variants. For unrestricted DiveRS versions, the number of unused reviewers from the original PC also lies around the value produced by PR4All. Of all methods, it is only DiveRS that prevents the generation of assignments with dependencies between reviewers in sets.




Figure 2: Reviewer assignment results using manual evaluation, displaying average diversity (x-axis) against the number of suitable reviewers (y-axis). The error bars correspond to the standard deviation per method. Results are reported on the two datasets combined.

For Γ 20 with $D_{\theta=.25}$ we found four papers as well as four original reviewers which were out of scope of the conference. For Γ 19 with $D_{\theta=.15}$ we found two original reviewers which were out of scope of the conference.

6.1.2 Manual Evaluation. We set up an online questionnaire where the two respective groups of PC chairs assessed the suitability of reviewer sets for ten randomly drawn submissions for their conferences. We presented them with four reviewer sets:⁵ the real assignment as well as three automatic assignments produced by PR4All, $D_{\theta=0}$, and $D_{\theta>0}$. For each assignment, PC chairs indicated the set’s suitability on a four-point scale (*no reviewers are suitable, two reviewers need to be replaced, one reviewer needs to be replaced, suitable assignment*) and justified their decision in a free-text field.

Figure 2 shows the average diversity against the number of suitable reviewers, for the two datasets combined. Both $D_{\theta=0}$ and $D_{\theta>0}$ produce reviewer sets with fewer suitable reviewers than the real assignment and PR4All—on the other hand, they produce much more diverse assignments. We observed low agreement between PC chairs when asked about the suitability of reviewer sets, as reflected in the standard deviations. It suggests that there are additional factors that may need to be considered in the reviewer assignment task; the free text comments, however, did not allow us to identify any common patterns.

6.1.3 Summary of Findings. DiveRS achieves fairness values which are comparable to those achieved by PR4All, without specifically focusing on this aspect of the problem. Additionally, our approach introduces more options to control reviewer load and to ensure the independence of reviewers. The resulting diversity values for DiveRS are much better than those of the real assignments or PR4All.

In our experiments, we found that there is a high probability of not assigning papers to all reviewers from the initial PC. Some members might have been included in a PC solely due to their reputation, not because of current interests or expertise in the fields of the submissions [2]. Unlike other methods, DiveRS offers the possibility of enforcing the involvement of all (fitting) PC members.

Manual evaluation showed the difficulty of objectively assessing the suitability of reviewer assignments, as we observed a high degree of disagreement between PC chairs. A comparison of diversity

⁵If sets produced from different methods are identical, we only depict it once.

Table 2: Reviewer suggestion results, listing average values for relevance of explanation (r), confidence (f), usefulness (u), convincingness (c), as well as suggestion ranking (NDCG) per dataset. Usefulness and convincingness are further subdivided (in parentheses) to cases with relevance below 3 ($u_{<}$, $c_{<}$) and above 3 ($u_{>}$, $c_{>}$).

d.set	r	f	u ($u_{<}/u_{>}$)	c ($c_{<}/c_{>}$)	NDCG
I'19	2.22	4.06	2.56 (2.15/3.67)	2.06 (1.69/3)	.7967
I'20	2.65	3.65	2.2 (1.89/2.17)	2.25 (1.56/3.17)	.9105

against the number of suitable reviewers revealed DiveRS' tendency to sacrifice some suitability in order to achieve high diversity.

6.2 Part 2: Reviewer Suggestion

In the second part of the evaluation, we measure the quality of reviewer suggestions for inclusion in the PC; this corresponds to our main routine (§4.2.3). We consider up to ten reviewer candidates suggested by DiveRS (6 for I'19 and 10 for I'20⁶). PC chairs are given a list of reviewers that could be invited. Each candidate is presented by their name, link to their DBLP profile, their main diversity attributes (professional background, location, seniority) as well as an explanation why they would be useful for an exemplary submission (e.g., *non-academia and academia background, topically fitting*). Additionally, other submissions in which the suggested candidate could help are listed. PC chairs are then asked to rate the relevance of the suggestion, their confidence in their assessment, as well as the usefulness of the explanation and how convincing it is on a Likert scale from 1 (not at all) to 5 (very).

The PC chairs' agreement on the relevance of suggestions is low for both datasets, which leads us to believe that this task is also very difficult to evaluate. The average values for assessed quality dimensions of suggested reviewer candidates are listed in Table 2. In general, relevancy for suggested reviewers is low, usefulness and convincingness of explanations increase drastically if only relevant (relevancy>3) are considered. We also evaluate suggestions as a ranked list in terms of NDCG, and observe high scores, especially for I'20. This can be interpreted as our method's ability to estimate the confidence of the recommendations and rank them accordingly. Our results hint at difficulties in suggestions' quality assessment, which should be investigated further to make findings more conclusive.

7 CONCLUSION

In this paper we introduced the novel reviewer coverage problem and proposed DiveRS, a flow-based reviewer assignment and PC member suggestion approach to solve it. DiveRS constructs diverse and fair reviewer set assignments for submissions and also suggests new reviewer candidates for inclusion in the PC. Our evaluation on two real world datasets showed DiveRS' superior diversity compared to both real assignments and the current state-of-the-art. Our experiments also highlighted the inherent difficulties of the reviewer assignment task, as evidenced by the low inter-annotator agreement between former PC chairs.

Future work could include utilising bidding information, when available, to identify papers with insufficient coverage. Requiring

junior reviewers to be part of each reviewer set may be desirable at times. Also, candidate suggestions may be subjected to stricter requirements, e.g., they should be able to review multiple submissions or not be considered at all. Additionally, creating a reusable dataset for reviewer suggestion will be a challenge in itself. Finally, there are further gains to be made by employing more advanced methods for determining the similarity between reviewers and submissions.

ACKNOWLEDGMENTS

We want to thank the conference general chairs of ICTIR'19 and ICTIR'20 for kindly providing us with the conference data, as well as the PC chairs for evaluating our automatic recommendations.

REFERENCES

- [1] Ken Birman and Fred B. Schneider. 2009. Viewpoint - Program committee overload in systems. *Commun. ACM* 52, 5 (2009), 34–37.
- [2] Georgeta Bordea, Toine Bogers, and Paul Buitelaar. 2013. Benchmarking domain-specific expert search using workshop program committees. In *CompSci@CIKM '13*. 19–24.
- [3] Laurent Charlin, Richard S. Zemel, and Craig Boutilier. 2011. A Framework for Optimizing Paper Matching. In *UAI '11*. 86–95.
- [4] Don Conry, Yehuda Koren, and Naren Ramakrishnan. 2009. Recommender systems for the conference paper assignment problem. In *RecSys '09*. 357–360.
- [5] Fred Douglass. 2008. Best Practices for the Care and Feeding of a Program Committee, and Other Thoughts on Conference Organization. In *WOWCS '08*.
- [6] Shuguang Han, Jiepu Jiang, Zhen Yue, and Daqing He. 2013. Recommending program committee candidates for academic conferences. In *CompSci@CIKM '13*. 1–6.
- [7] Steven Jeemen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang. 2020. Mitigating Manipulation in Peer Review via Randomized Reviewer Assignments. In *NeurIPS '20*.
- [8] Jian Jin, Qian Geng, Haikun Mou, and Chong Chen. 2019. Author-Subject-Topic model for reviewer recommendation. *J. Inf. Sci.* 45, 4 (2019).
- [9] Noriaki Kawamae. 2010. Author interest topic model. In *SIGIR '10*. 887–888.
- [10] Ari Kobren, Barna Saha, and Andrew McCallum. 2019. Paper Matching with Local Fairness Constraints. In *SIGKDD '19*. 1247–1257.
- [11] Ngai Meng Kou, Leong Hou U, Nikos Mamoulis, and Zhiguo Gong. 2015. Weighted Coverage based Reviewer Assignment. In *SIGMOD '15*. 2031–2046.
- [12] Michael Ley. 2009. DBLP - Some Lessons Learned. *Proc. VLDB Endow.* 2, 2 (2009), 1493–1500.
- [13] Xiang Liu, Torsten Suel, and Nasir D. Memon. 2014. A robust model for paper reviewer assignment. In *RecSys '14*. 25–32.
- [14] Cheng Long, Raymond Chi-Wing Wong, Yu Peng, and Liangliang Ye. 2013. On Good and Fair Paper-Reviewer Assignment. In *ICDM '13*. 1145–1150.
- [15] Haikun Mou, Qian Geng, Jian Jin, and Chong Chen. 2015. An Author Subject Topic Model for Expert Recommendation. In *AIRS '15*. 83–95.
- [16] Fabio Pacheco, Igor Wiese, Bruno Cartaxo, Igor Steinmacher, and Gustavo Pinto. 2020. Analyzing the evolution and diversity of SBES Program Committee. *CoRR* (2020).
- [17] Simon Price and Peter A. Flach. 2017. Computational support for academic peer review: a perspective from artificial intelligence. *Commun. ACM* 60, 3 (2017), 70–79.
- [18] Michal Rosen-Zvi, Thomas L. Griffiths, Mark Steyvers, and Padhraic Smyth. 2004. The Author-Topic Model for Authors and Documents. In *UAI '04*. 487–494.
- [19] Vyas Sekar. 2016. EZ-PC: program committee selection made easy. *Comput. Commun. Rev.* 46, 3 (2016), 6:1–6:3.
- [20] Anna Severin and Joanna Chataway. 2021. Overburdening of peer reviewers. A multi-disciplinary and multi-stakeholder perspective on causes, effects and potential policy implications. *bioRxiv* (2021).
- [21] Ivan Stelmakh, Nihar B. Shah, and Aarti Singh. 2018. PeerReview4All: Fair and Accurate Reviewer Assignment in Peer Review. *CoRR* (2018).
- [22] Wenbin Tang, Jie Tang, and Chenhao Tan. 2010. Expertise Matching via Constraint-Based Optimization. In *WT '10*. 34–41.
- [23] Yuancheng Tu, Nikhil Johri, Dan Roth, and Julia Hockenmaier. 2010. Citation Author Topic Model in Expert Search. In *COLING '10*. 1265–1273.
- [24] Bogdan Vasilescu, Alexander Serebrenik, Tom Mens, Mark G. J. van den Brand, and Ekaterina Pek. 2014. How healthy are software engineering conferences? *Sci. Comput. Program.* 89 (2014), 251–272.
- [25] Thanasis Vergoulis, Serafeim Chatzopoulos, Theodore Dalamagas, and Christos Tryfonopoulos. 2020. VeTo: Expert Set Expansion in Academia. In *TPDL '20*. 48–61.

⁶DiveRS introduces different numbers of reviewers based on the dataset as well as θ .