

000 PAPER COPILOT: TRACKING THE EVOLUTION OF PEER 001 REVIEW IN AI CONFERENCES 002

003 **Anonymous authors**

004 Paper under double-blind review

005 ABSTRACT

006 The rapid growth of AI conferences is straining an already fragile peer-review
007 system, leading to heavy reviewer workloads, expertise mismatches, inconsistent
008 evaluation standards, superficial or templated reviews, and limited accountability
009 under compressed timelines. In response, conference organizers have introduced
010 new policies and interventions to preserve review standards. Yet these ad-hoc
011 changes often create further concerns and confusion about the review process,
012 leaving how papers are ultimately accepted—and how practices evolve across
013 years—largely opaque. We present PAPER COPILOT, a system that creates durable
014 digital archives of peer reviews across a wide range of computer-science venues, an
015 open dataset that enables researchers to study peer review at scale, and a large-scale
016 empirical analysis of ICLR reviews spanning multiple years. By releasing both the
017 infrastructure and the dataset, PAPER COPILOT supports reproducible research on
018 the evolution of peer review. We hope these resources help the community track
019 changes, diagnose failure modes, and inform evidence-based improvements toward
020 a more robust, transparent, and reliable peer-review system.

021 1 INTRODUCTION

022 The rapid growth of submissions to top-tier Artificial Intelligence (AI) and Machine Learning (ML)
023 conferences—now exceeding 10,000 per venue annually (Yang, 2025)—has placed unprecedented
024 pressure on peer review. To address issues of scale and transparency, many conferences have adopted
025 open or semi-open platforms like OpenReview. Yet practices remain inconsistent: some venues
026 publish reviews and scores, while others keep them private. This inconsistency has sparked ongoing
027 debate around fairness, accountability, and the effectiveness of different review models (Cortes &
028 Lawrence, 2021).

029 Meanwhile, the nature of review data itself has evolved. Whereas reviews were initially limited
030 to a single numerical rating, nowadays reviews include multiple dimensions such as soundness,
031 technical correctness, and presentation quality (Beygelzimer et al., 2021). Since the rise of Open-
032 Review (Soergel et al., 2013) and especially following the widespread adoption of formal rebuttal
033 phases, the need to analyze score dynamics over time has grown (Yang, 2025). Yet community
034 discourse remains scattered—spread across Twitter, Reddit, Zhihu, Xiaohongshu, and other social
035 platforms in fragmented, venue-specific threads. As submission volumes and review complexity
036 continue to increase, the absence of structured tools that unify and visualize the review data has
037 become a bottleneck for both transparency and timely author decision-making.

038 From the standpoint of conference organizers, internal dashboards provide visibility into score trends
039 and acceptance distributions (Cortes & Lawrence, 2021). But these views are rarely made accessible
040 to authors. During the critical post-review phase—often only one or two weeks long—authors are
041 expected to analyze reviewer feedback, prepare rebuttals, and potentially decide whether to withdraw
042 or revise their submissions. With scattered resources, authors must resort to scraping statistics (Sun,
043 2020) or aggregating posts from scattered social media. This inefficiency hinders timely decision-
044 making and may even lead to missed opportunities for meaningful rebuttal or clarification. A system
045 that actively collects review statistics, visualizes score dynamics, and contextualizes submissions
046 relative to their peers would therefore offer immediate utility to thousands of authors.

047 In addition, while the OpenReview ecosystem has grown, many conferences still use closed-form
048 review systems and opt not to release reviews or scores publicly. Motivations for withholding

054 reviews may include protection of intellectual content, reviewer anonymity, or historical precedent.
 055 However, survey studies of peer review in ML have documented that such opacity exacerbates
 056 known problems: inconsistent reviewing standards, lack of calibration across reviewers, and limited
 057 accountability for low-quality or biased reviews (Shah, 2022). These shortcomings directly affect
 058 fairness and community trust, and they compound when authors have no visibility into aggregate
 059 statistics. A practical workaround is to allow voluntary, anonymized community submissions. If
 060 designed carefully, such a system can extract value from partially open data while respecting privacy
 061 and consent.

062 Beyond the review process itself, the AI / ML community lacks robust infrastructure for tracking who
 063 is shaping the field over time. Traditional platforms like Google Scholar and DBLP focus primarily on
 064 paper-level or author-level metadata—citations, references, and publication profiles (Scholar; DBLP).
 065 While Google Scholar does display author affiliations (and, by extension, country information), these
 066 systems were not designed to support temporal analysis or dynamic tracking across institutions or
 067 regions. Their scope remains limited, offering little visibility into broader affiliation-, institution-, or
 068 country-level dynamics. Tracking such information has historically been left to organizations due
 069 to the technical and logistical challenges involved, as seen in efforts like CSRankings (Berger) or
 070 university rankings.

071 As AI becomes increasingly high-profile and competitive, attention is increasingly concentrated
 072 around specific venue cycles. This shift highlights a growing need to understand who is actively
 073 driving progress over shorter timeframes: which institutions are rising, who remains active in the field,
 074 and what geographic regions are gaining influence. Yet no existing academic tool provides this kind
 075 of dynamic, multi-scale view. Ranking systems often update annually and draw from non-transparent
 076 data sources, making them ill-suited to capture the rapid and evolving dynamics of the field.

077 To address these needs and challenges, we introduce **Paper Copilot**, a system with open dataset for
 078 tracking peer review dynamics and talent trajectories across AI/ML. Our contributions include:

- 079 • **System for tracking AI/ML progress.** A venue-configurable pipeline (Fig. 1) that unifies
 080 multi-source inputs into standardized, versioned paperlists and powers interactive, multi-
 081 granularity analytics (venue/institution/country) for longitudinal progress tracking.
- 082 • **Scalable peer-review archive.** A unified archive built from open, semi-open, and opt-in
 083 community data with temporal snapshots, capturing multi-dimensional review metadata and
 084 daily dynamics (ICLR 2024/2025), to preserve and analyze cross-venue review evolution at
 085 scale.
- 086 • **Findings & ethical considerations.** Empirical results show a 2025 shift toward sharper,
 087 score-driven tiering under volume pressure and characteristic rebuttal-phase dynamics (score
 088 shifts, consensus evolution); we articulate safeguards on sourcing/consent, privacy, misuse,
 089 and bias.

091 2 RELATED WORKS

092 2.1 PEER REVIEW DATASETS AND ANALYSIS

093 There is a growing body of work on creating datasets to study academic peer review. PeerRead (Kang
 094 et al., 2018), collected 14.7K paper drafts with reviews and decisions from NLP / ML venues including
 095 ACL, NeurIPS and ICLR, enabling research on review–decision alignment and NLP applications like
 096 review score prediction. More recently, MOPRD (Lin et al., 2023) introduced a multidisciplinary
 097 open peer review dataset spanning several journals and computer science conferences, capturing
 098 the full reviewing process (including review comments, rebuttals, meta-reviews and final decisions)
 099 across domains. Other efforts have targeted specific aspects: for example, ORB (Szumega et al.,
 100 2023) curated a list of more than 36,000 scientific papers with their more than 89,000 reviews and
 101 final decisions in high energy physics.

102 Bharti et al. (2022) leveraged an NLP approach to estimate reviewer confidence from review text,
 103 illustrating analysis that can be done when reviews are available. Peer Review Analysis (Ghosal et al.,
 104 2022) provided 1,199 review reports labeled at the sentence level for aspects like critique, suggestion,
 105 and sentiment, serving as a benchmark for automated review assessment. Ribeiro et al. (2021) used
 106 a private dataset of two non-disclosed conferences to perform RSP, PDP, and sentiment analysis to

108 extract review polarities. In general, these datasets have been invaluable in enabling studies on review
 109 quality, bias, and consistency. However, each is often limited to a particular venue or a snapshot in
 110 time. We advance beyond prior datasets by unifying multiple venues over a long temporal window,
 111 thereby supporting cross-conference and longitudinal analyses (e.g., how review score distributions
 112 shift over years, or how an author’s successive submissions fare over time). It also introduces the
 113 notion of tracking review process “dynamics,” since our data include timestamps and sequences
 114 (allowing one to reconstruct the timeline of reviews and discussions for open-review conferences).
 115 This complements earlier one-off studies that examined review consistency and randomness in single
 116 years (Beygelzimer et al., 2021; Cortes & Lawrence, 2021), by providing a broader data foundation
 117 to study such effects in the aggregate.

118 **2.2 PEER REVIEW TRACKING SYSTEMS**

119 Our work is related to platforms that manage or expose the peer review process. OpenReview.net is
 120 the most prominent system enabling open peer review for many ML conferences (ICLR, NeurIPS,
 121 UAI, etc.), providing public APIs to fetch submissions, reviews, and comments (OpenReview, 2024).
 122 We extensively leverages OpenReview as a data source, but whereas OpenReview is focused on
 123 facilitating the review process for active conferences, our system is designed to track and analyze
 124 review outcomes across conferences and years. In some sense, Paper Copilot serves as a meta-layer
 125 on top of systems like OpenReview and other open-access portals and proceedings (Foundation, 2024;
 126 Anthology), aggregating their outputs for comparative analysis. A few individual conferences have
 127 released summary statistics (e.g., acceptance rates, score distributions) in blog posts or slides, but
 128 these are ad-hoc and not standardized (Beygelzimer et al., 2021). To our knowledge, there was no
 129 public web resource before Paper Copilot that allowed users to conveniently browse and compare
 130 peer review results (scores, decisions, etc.) for a wide range of AI / ML venues in one place. Our
 131 system fills this gap by acting as a centralized peer review statistics dashboard for the community. It
 132 is somewhat analogous to “CS Conference Stats” websites that track acceptance rates (Berger), but
 133 we go further by including detailed review metrics and providing the underlying dataset.

134
 135 **2.3 AUTHOR PROFILING AND METASCIENCE**

136 Studying author career trajectories and publication patterns has been a long-standing interest in
 137 bibliometrics (Fan et al., 2024; De Frutos Belizon et al., 2025). Large-scale academic databases and
 138 knowledge graphs – such as Google Scholar, Microsoft Academic Graph, Semantic Scholar’s Open
 139 Research Corpus, and AMiner – have enabled analyses of collaboration networks, citation trajectories,
 140 and researcher mobility (Scholar). For example, AMiner (Tang et al., 2008), an academic social
 141 network mining system that compiled tens of thousands of author profiles and their affiliations over
 142 time, allowing for queries on career moves and productivity. A recent dataset from AMiner contains
 143 over 57,000 scholars’ career paths (with 42,230 affiliations) for studying academic job transitions.
 144 These tools, however, focus on publication and citation data; they generally do not incorporate the
 145 peer review stage. We brings a new dimension by linking authors to the review process outcomes
 146 of their submissions. This creates opportunities for “talent trajectory” studies in AI / ML – for
 147 instance, one could track how an early-career researcher improves their paper quality (as reflected in
 148 review scores) over successive submissions, or examine whether consistent high reviewers’ ratings
 149 correlate with future citation impact (Vinkenburg et al., 2020). Our dataset, combined with external
 150 bibliometric data, could also help identify patterns such as mentorship lineages (e.g., how students
 151 of certain groups perform in reviews) or geographic trends in research feedback (Guevara et al.,
 152 2016). In the broader context of metascience, our work contributes an open resource to empirically
 153 study questions about peer review and researcher development. This aligns with recent calls for
 154 transparency and data-driven policy in peer review (Bianchi & Squazzoni, 2022). By making our
 155 data easily accessible, we hope to facilitate further meta-research on how review processes influence
 156 and reflect the growth of the AI / ML research community.

157
 158 **3 SYSTEM**

159 Paper Copilot is a modular system for large-scale collection and presentation of peer review dynamics
 160 and talent trajectories. Figure 1 outlines its architecture: venue configurations define scalable sources,
 161 processed through assigners, worker pools, and parallel bots to generate standardized paperlists.

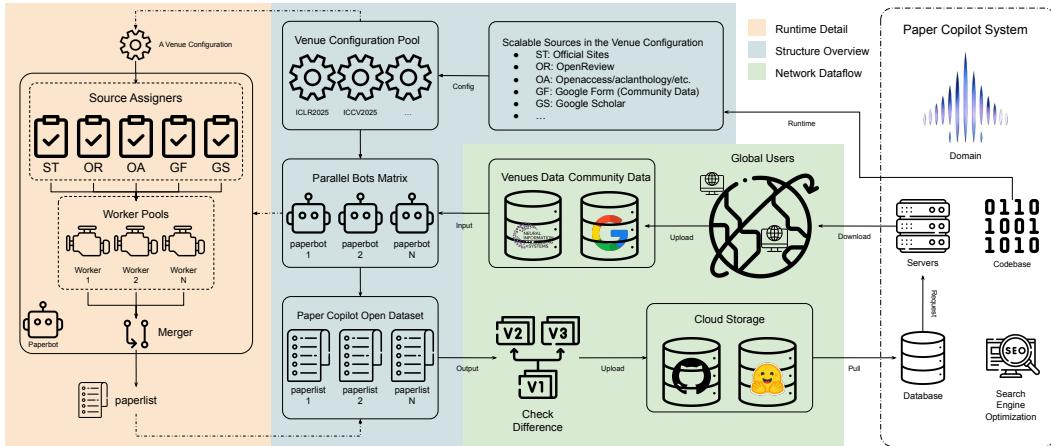


Figure 1: Paper Copilot System Architecture

Versioned datasets are validated, uploaded to cloud storage, and served globally through optimized backend, database, and frontend components. The following sections detail data collection, system backend/frontend, and maintenance.

Data Collection We use a hybrid data ingestion pipeline to accommodate diverse conference policies, including: 1) *OpenReview API*: For open-review venues (e.g., ICLR, NeurIPS), scheduled scripts pull submission metadata, review scores, confidences, and discussion threads. We store timestamped snapshots to track score changes (e.g., pre-/post-rebuttal). 2) *Website Scraping*: For venues without APIs (e.g., CVPR, AAAI), lightweight scrapers extract accepted papers, authors, and other publicly available metadata from official websites or proceedings. 3) *Community Contributions*: For closed-review venues, authors voluntarily submit reviews via a structured form, yielding 6,584 valid reviews per submission to date.

All data flows through a cleaning and normalization module that standardizes fields, deduplicates entries, and prepares them for integration. Adding new venues requires minimal configuration.

Backend and Frontend The backend is built on a LAMP stack (Linux, Apache, MySQL, PHP) and deployed via a cPanel-managed hosting environment, which simplifies cron scheduling, backups, and database management. Conference data—submissions, reviews, authorship—is stored in a MySQL relational database with a normalized schema linking papers, reviews, and authors. Pre-computed aggregates (e.g., average score, acceptance rate per venue/year) are cached to support fast retrieval for analytics and display.

The user-facing website is built with WordPress, extended by custom PHP and JavaScript for dynamic rendering. All pages support filtering by score range, decision status, or specific metadata fields, and provide links to external resources such as OpenReview, official proceedings, or ArXiv when available. Each conference currently has two primary views: 1) *Statistics*: Interactive charts display trends in paper counts, acceptance rates, score distributions, and reviewer score rankings. 2) *Paper list*: Tabulated entries show paper titles, primary areas, authors, affiliations, departments, countries, citations, and acceptance types (e.g., oral/poster).

Beyond individual venues, the system supports cross-venue, longitudinal analysis at the institutional level. Users can explore hierarchical structures—tracking contributions over time from institutions to departments, laboratories, and regions. This enables global and fine-grained insights into talent flows, geographic shifts, and institutional participation across the field. For more advanced exploration, we provide a companion search tool that can run locally. Users can query by score thresholds, keywords, or metadata fields, operating on either the live database or JSON exports.

Maintenance and Updates The entire Paper Copilot platform—from backend to frontend—is fully automated and designed for scalability, reliability, and long-term meta-research with minimal manual intervention. Cron jobs manage regular updates to review dynamics; scraper failures trigger alerts; and the database is routinely backed up to ensure data integrity. Static pages are cached to reduce server load, and the modular pipeline allows easy adaptation as conference policies evolve. Developed and maintained at low cost, the system’s modular architecture enables rapid scaling to new venues

216 with minimal overhead. All processed datasets are publicly available, as detailed in Section 4.2,
 217 supporting transparency, reproducibility, and community-driven enhancements to features such as
 218 search and visualization.

219 4 DATASET

220 The Paper Copilot open dataset spans decades of peer review data and accepted paper metadata across
 221 dozens of AI / ML conferences and continues to grow. In this section, we detail the dataset’s contents,
 222 structure, and coverage. We describe how temporal review profiles are represented, the coverage of
 223 authors and venues, and provide examples of insights one can derive. We also explain the format in
 224 which we release the data and the visualizations included.

225 4.1 PEER REVIEW DYNAMICS

226 **Review Dimensions** As AI / ML research expands, review criteria have grown in both number and
 227 complexity, varying across venues and years. In addition to core metrics like rating and confidence,
 228 many conferences now assess dimensions such as soundness, correctness, novelty, contribution,
 229 presentation, etc. These evolving schemas reflect a shift toward more granular peer evaluation.

230 **Reviewer-Author Discussion Score Dynamics** For venues with fully open review processes (e.g.,
 231 ICLR), we leverage the public API to track the temporal evolution of review scores, confidence,
 232 and comments throughout the discussion and rebuttal phases. This allows us to construct reviewer
 233 profiles over time and analyze how feedback shifts during the review cycle. By timestamping each
 234 review snapshot, our system captures pre- and post-rebuttal score changes at the individual reviewer
 235 level, enabling fine-grained analysis of reviewer behavior and consistency. Figure 3 illustrates these
 236 dynamics using ICLR 2025 public data. Although ICLR reviews are public, historical snapshots are
 237 not preserved on the official platform—older versions are overwritten during the discussion phase.
 238 Paper Copilot continuously archives these updates in real time, providing what is, to the best of our
 239 knowledge, the **only publicly accessible archive of review score dynamics available anywhere on**
 240 **the internet.**

241 For closed-review venues that do not disclose review revisions (e.g., CVPR, ICCV, ICML), we invite
 242 authors to voluntarily submit their initial and final review scores. This opt-in process enables partial
 243 reconstruction of rebuttal impact, even when the official platform lacks transparency. By combining
 244 automated tracking and community-contributed data, our system enables venue-level comparisons of
 245 rebuttal effectiveness and decision volatility across time.

246 **Closed Venue Review Disclosure** Due to the closed nature of most conference reviews, public
 247 access to detailed review data remains limited. As part of our community engagement, we launched
 248 an opt-in questionnaire alongside our score collection process to assess author willingness to anonym-
 249 ously share their reviews. This study was conducted across major venues, including CVPR, ICML,
 250 ICCV, and ACL, and will continue with NeurIPS 2025.

251 To characterize community attitudes toward transparency, we conducted a user study across four
 252 major conferences in 2025. In total, we collected 1,860 responses, of which 1,115 authors (59.9%)
 253 consented to publicly release their anonymized review scores. Consent rates were broadly consistent
 254 across venues: 53.5% at CVPR (191 of 357), 60.7% at ICML (628 of 1,034), 59.4% at ICCV (151
 255 of 254), and 67.4% at ACL (145 of 215). These results indicate strong support—roughly 60%
 256 overall—for releasing anonymized review data. We plan to continue this initiative across future years
 257 and venues to build a longitudinal dataset that can inform empirical studies of peer review models.

258 4.2 DATASET FORMAT AND RELEASE

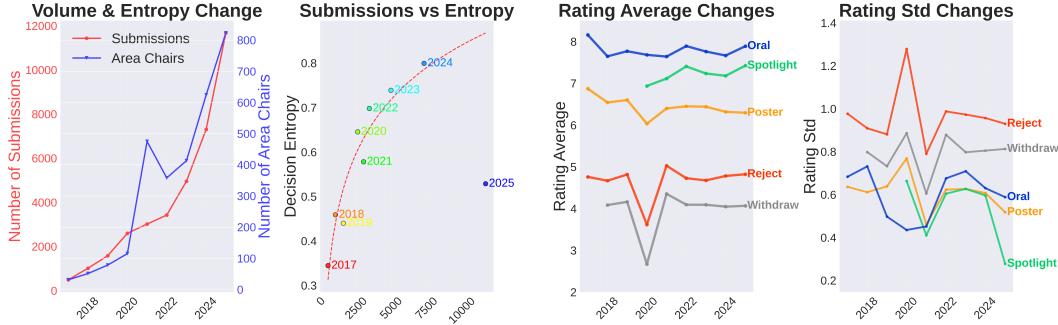
259 To ensure usability and reproducibility, we release the dataset in structured JSON format, with one
 260 record per paper containing its metadata, review scores, timestamps, and author/affiliation information.
 261 Each paper record includes over 30 fields covering metadata, authorship, review scores, rebuttal
 262 dynamics, and final decisions. More details of data acquisition can be found in the supplementary
 263 material. The main dataset is available on GitHub A separate repository, GitHub, hosts the temporal

270 data with parsing code for tracking review dynamics over time, due to size constraints. We are also
 271 collecting ICLR 2026 review data, which will be integrated into the archive in future updates if the
 272 conference continues to adopt a fully open review model.

274 5 THE EVOLUTION OF PEER REVIEW IN AI CONFERENCES.

277 In this section, we analyze the evolution of peer review in ICLR, a frontier AI conference. Our
 278 analysis aims at two key questions: (1) how has the conference decision-making mechanism evolved
 279 across years and (2) what are the discussion dynamics of ICLR during the rebuttal period?

280 5.1 THE EVOLUTION OF CONFERENCE DECISION-MAKING MECHANISMS



293 Figure 2: ICLR review dynamics from 2017 to 2025. Despite stable score distributions, decision
 294 entropy rises with scale—until 2025, where it drops sharply. This deviation suggests Area Chairs
 295 play a more decisive role, increasingly relying on mean scores under high submission pressure.

297 Over the past decade, ICLR has experienced explosive growth—from just 490 submissions in 2017
 298 to 11,672 in 2025 (Figure 2, leftmost panel). To manage this scale, the number of Area Chairs
 299 (ACs) expanded proportionally, rising from 31 to 823 over the same period. This rapid expansion
 300 underscores the systemic pressure imposed by volume growth.

302 When analyzing the **final decisions** made by ACs, we quantify uncertainty using **decision entropy**,
 303 which measures unpredictability in tier assignment given a paper’s mean score. Formally, for year t
 304 and score bin b :

$$305 H_{t,b} = - \sum_{s \in \{\text{Reject, Poster, Spotlight, Oral}\}} p_{t,b,s} \log p_{t,b,s}, \quad \bar{H}_t = \sum_b w_{t,b} H_{t,b},$$

308 where $p_{t,b,s}$ is the empirical probability of status s for papers with binned mean score, and $w_{t,b}$ are
 309 bin-level weights. Higher \bar{H}_t indicates greater uncertainty; lower values indicate sharper, score-driven
 310 decisions.

311 As shown in Figure 2 (center-left), \bar{H}_t typically increases with submission volume X_t , and is well
 312 approximated by a logarithmic scaling:

$$313 \bar{H}_t \approx a \log X_t + b.$$

315 This trend aligns with an ordered-logit model for score-to-status mapping:

$$316 P(\text{status} = s \mid x) = \sigma(\tau_s - \kappa_t x) - \sigma(\tau_{s-1} - \kappa_t x),$$

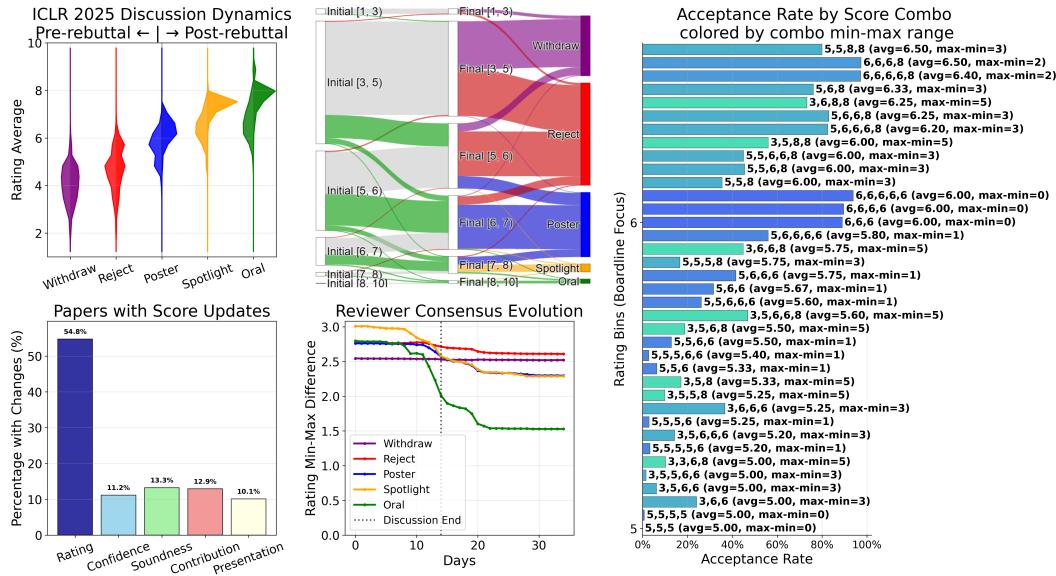
318 where σ is the logistic CDF and κ_t reflects decision sensitivity to mean score in year t . As X_t grows,
 319 κ_t typically softens, yielding higher entropy.

320 **Yet ICLR 2025 is a clear outlier.** Despite the largest submission volume ($X_{2025} = 11,672$),
 321 \bar{H}_{2025} fell below the fitted log trend. This implies a stronger κ_{2025} , i.e., the venue relied more
 322 deterministically on average scores, reducing uncertainty rather than increasing it. The residual,

$$323 \text{resid}_{2025} = \bar{H}_{2025} - (a \log X_{2025} + b),$$

324 is strongly negative, showing that scale alone cannot explain the drop.
 325
 326 Turning to score allocation (Figure 2, center-right and rightmost panels), **average ratings** have grown
 327 more tier-separated, with Spotlights converging toward Orals. Meanwhile, **rating variance** rises
 328 for rejected/withdrawn papers but drops for accepted ones, especially Spotlights. This shows ACs
 329 increasingly favor high and stable scores, sharpening thresholds and limiting borderline flexibility.
 330
 331 Together, these results indicate a structural change in 2025: ACs assumed a more decisive role, enforc-
 332 ing sharper score-dependent rules. The system shifted from probabilistic tiering to near-deterministic
 333 mappings, improving efficiency but potentially narrowing the role of reviewer justifications, rebuttals,
 334 or confidence.

335 5.2 DISCUSSION DYNAMICS



355
 356 Figure 3: **Rebuttal dynamics in ICLR 2025.** (top-left) Ridge plots of pre- vs. post-rebuttal ratings
 357 across final decision tiers. (top-mid) Sankey diagram tracking score trajectories from initial rating bins
 358 to final bins and final statuses. (lower-left) Distribution of score updates across review dimensions.
 359 (lower-mid) Reviewer consensus evolution measured by min–max rating difference. (right) Accep-
 360 tance rates conditioned on reviewer score combinations with frequency cutoff=30.

361 To complement the volume and entropy analysis, we examine how the rebuttal phase shaped reviewer
 362 behavior and final decisions in ICLR 2025 (Figure 3).

363
 364 **Score shift before and after rebuttal** The ridge plot (top-left) shows that score shifts after rebuttal
 365 were most pronounced in higher tiers. Spotlight and Oral papers experienced clear upward movement,
 366 whereas Withdrawn and Rejected papers showed little change. The Sankey diagram (top-right) further
 367 traces the evolution of scores from initial assignments to post-rebuttal updates and final status. Here,
 368 green flows denote upward shifts in mean score and red flows denote downward shifts. The dominant
 369 pattern is that many borderline and mid-tier papers received upward adjustments, while downward
 370 movements remain less common but non-negligible.

371
 372 **Reviewer score change patterns and consensus** At a finer granularity, score updates were not
 373 uniform across review dimensions (bottom-left). Over half of papers (54.8%) saw changes in the
 374 overall rating, whereas dimensions such as confidence, soundness, contribution, and presentation
 375 shifted in only \sim 10–13% of cases. This suggests that rebuttals primarily influenced holistic ratings
 376 rather than specific dimension scores.

377 Reviewer consensus dynamics (bottom-center) show that rating min–max differences consistently
 378 narrowed as discussions unfolded. Notably, consensus levels dropped sharply at the onset of the

378 discussion phase (marked by the vertical dashed line), reflecting the divergence that arises when
 379 reviewers actively debate rebuttal content. Toward the end of the process, consensus recovered, with
 380 Oral papers converging most strongly by the decision deadline, while Rejected papers retained higher
 381 levels of disagreement. This indicates that high-stakes papers underwent deeper discussion and
 382 alignment, while weaker papers received less effort in consensus-building.
 383

384 **Acceptance rate and score patterns** Finally, we inspect acceptance rates at the borderline using
 385 reviewer score combinations (bottom-right). After filtering for frequency (cutoff = 30), we find
 386 a notable asymmetry: when the average score is slightly above the borderline, lower min–max
 387 ranges correlate with higher acceptance rates; conversely, when the average is slightly below, higher
 388 min–max ranges correlate with acceptance. A plausible explanation is that below the borderline,
 389 a single strongly positive reviewer can sway the outcome, while above the borderline, persistent
 390 disagreement often signals unresolved concerns leading to rejection. While compelling, this empirical
 391 pattern warrants further validation in future analyses.

392 Taken together, these findings highlight the dual role of rebuttals: amplifying score changes for
 393 borderline cases and driving consensus formation for stronger papers. The corresponding plots for
 394 ICLR 2024 are provided in the appendix, enabling year-over-year comparison.
 395

396 6 ETHICAL CONSIDERATIONS

398 The development of a dataset tracking peer review and talent trajectories necessitates a careful and
 399 transparent approach to ethical considerations. We are committed to responsible data stewardship
 400 and have outlined our mitigation strategies for key ethical challenges below.
 401

402 6.1 DATA SOURCING, LICENSING, AND CONSENT

404 One potential concern revolves around the methods of data collection and adherence to licensing
 405 and Terms of Service (ToS). Our data is aggregated from three distinct sources, each with a different
 406 consent and licensing model. 1) **Public APIs and Data Sources:** A significant portion of our data is
 407 sourced from platforms with explicit data-sharing policies, such as OpenReview, which provides a
 408 public API for accessing review data. Data from such sources is collected in accordance with their
 409 established terms. 2) **Web Scraping:** Automated scraping can violate the ToS of some platforms,
 410 such as Google Scholar. To ensure full compliance, we have audited our data collection pipelines
 411 and ceased all automated data collection that conflicts with site ‘robots.txt’ files or explicit ToS. 3)
 412 **Community-Contributed Data:** For closed-review venues, we rely on voluntary, opt-in submissions
 413 from authors. Each author who contributes data does so via a form that includes explicit consent
 414 questions. Authors can choose whether their submission contributes only to anonymized, aggregate
 415 statistics or if their anonymized scores can be displayed in detailed tables showing individualized
 416 statistics.

417 6.2 PRIVACY, ANONYMITY, AND RE-IDENTIFICATION RISK

419 Another concern, potentially related to data licensing, is about protecting the identities of both authors
 420 and reviewers.

421 We do not collect, store, or attempt to de-anonymize reviewer identities. The vast majority of review
 422 data comes from platforms like OpenReview, where reviewer identities are already anonymized
 423 by default. There are potentially concerns that author trajectories, especially when combined with
 424 public timelines or multiple affiliations, could create re-identification risks. To mitigate this, our
 425 primary analysis focuses on **large-scale, aggregated trends at the institutional and geographic**
 426 **levels rather than individual-level data.** In addition, we are committed to implementing privacy-
 427 preserving measures, such as applying differential privacy to aggregate statistics.

429 6.3 POTENTIAL FOR MISUSE AND DUAL-USE RISKS

431 The dataset’s potential for misuse, particularly in hiring and evaluation, is a significant “dual-use”
 432 risk. A worst-case scenario involves the data being used to unfairly judge job candidates, especially

432 early-career researchers. Our stated goal is to increase transparency in the peer review process and
 433 empower researchers with a broader view of the academic landscape, not to create a performance
 434 ranking tool. To guard against misuse, we will take the following steps: 1) *Explicit Use Guidelines*:
 435 The Paper Copilot website will feature prominent guidelines strongly discouraging the use of its data
 436 for hiring, promotion, or other high-stakes evaluations of individuals. 2) *Focus on Aggregates*: The
 437 platform will prioritize the visualization of aggregated, large-scale trends over individual-level data.
 438 3) *Community Dialogue*: We are committed to engaging in an ongoing dialogue with the AI/ML
 439 community to monitor how the data is used and to develop stronger safeguards as needed.

440 6.4 DATA INTEGRITY, INACCURACY, AND BIAS

441 The dataset is subject to several sources of potential inaccuracy and bias.

442 **Sampling and Representation Bias.** For closed-review venues, the reliance on community sub-
 443 missions introduces a self-selection bias. While this is currently the only feasible method for tracking
 444 review dynamics at scale without official data access, we recognize its limitations. To rigorously quan-
 445 tify this sampling bias, we plan to conduct a comparative analysis between our community-submitted
 446 data and the official ground-truth data from NeurIPS 2025 once it is released.

447 **Data Inaccuracy.** Inaccurate data, whether from parsing errors or malicious submissions, could
 448 harm researchers' reputations. The LLM-based affiliation extraction has a non-zero error rate. We
 449 will mitigate this by: 1) Clearly labeling the source of all data (e.g., "Official API", "Community-
 450 Submitted"). 2) Providing transparent metrics on the known error rates of our extraction models. 3)
 451 Implementing validation checks for community submissions to flag anomalous entries.

452 **Demographic and Geographic Bias.** The dataset could contain demographic and geographic
 453 biases, posing a special risk to researchers from certain groups (e.g., by geography, gender, or race).
 454 This bias can manifest in several ways. For closed-review conferences, the dataset's reliance on
 455 voluntary community submissions may skew it toward certain regions or institutions that are more
 456 culturally inclined to share data. Furthermore, some research groups have internal policies that
 457 discourage sharing review data even when it is possible to opt-out, which could lead to the systematic
 458 underrepresentation of those groups in our dataset.

459 Such imbalances could lead to skewed analyses, where trends observed in overrepresented groups
 460 are incorrectly generalized to the entire research community, potentially disadvantaging those in
 461 underrepresented regions. To actively address this, we are committed to the following actions:

- 462 • **Transparency via Datasheets:** We will publish and maintain a datasheet that transparently
 463 documents the known geographic and institutional distributions in our dataset, highlighting
 464 potential skews so that researchers can interpret the data with proper context.
- 465 • **Regular Bias Audits:** We will perform regular audits to quantify these biases. Where
 466 possible, we will compare the distribution in our data against available ground-truth statistics
 467 from conference organizers to measure the extent of the disparity.
- 468 • **Platform-Level Context:** On our platform, any analyses derived from potentially biased
 469 data sources, such as community submissions, will be clearly marked with disclaimers that
 470 explain the data's limitations and advise against overgeneralization.

471 7 CONCLUSION

472 Paper Copilot provides a unified, large-scale platform for tracking peer review dynamics and talent
 473 trajectories across major AI / ML conferences. It enables comparative analysis of score trends,
 474 rebuttal effectiveness, and institutional influence, with datasets already adopted by both academia and
 475 industry. Key limitations include partial reliance on voluntary data contributions for closed-review
 476 venues and the need to further improve the accuracy of talent trajectory reconstruction. Future work
 477 will focus on expanding venue coverage and enhancing metadata quality for more comprehensive
 478 and precise analysis.

486 REFERENCES
487

488 ACL Anthology. Acl anthology. <https://aclanthology.org/>. Accessed: 2025-05-15.

489 Emery D. Berger. Csrankings: Computer science rankings. <https://csrankings.org/>.
490 Accessed: 2025-05-14.

491

492 Alina Beygelzimer, Yann Dauphin, Percy Liang, and Jennifer Wortman Vaughan. The neurips 2021
493 consistency experiment. *Neural Information Processing Systems blog post*, <https://blog.neurips.cc/2021/12/08/the-neurips-2021-consistency-experiment>, 2021.

494

495 Santosh Kumar Bharti, S Varadhanapathy, Rajeev Kumar Gupta, Prashant Kumar Shukla, Mohamed
496 Bouye, Simon Karanja Hingaa, and Amena Mahmoud. Text-based emotion recognition using deep
497 learning approach. *Computational Intelligence and Neuroscience*, 2022(1):2645381, 2022.

498

499 Federico Bianchi and Flaminio Squazzoni. Can transparency undermine peer review? a simulation
500 model of scientist behavior under open peer review. *Science and Public Policy*, 49(5):791–800,
501 2022.

502 Corinna Cortes and Neil D Lawrence. Inconsistency in conference peer review: Revisiting the 2014
503 neurips experiment. *arXiv preprint arXiv:2109.09774*, 2021.

504

505 DBLP. dblp computer science bibliography. monthly snapshot release of may 2025. <https://dblp.org/xml/release/dblp-2025-05-01.xml.gz>. Accessed: 2025-05-14.

506

507 Jesus De Frutos Belizon, Felix Guerrero Alba, and Gonzalo Sanchez Gardey. Exploring the academic
508 research career: A bibliometric and content analysis of the academic life cycle. *Innovative Higher
509 Education*, pp. 1–30, 2025.

510

511 Yangliu Fan, Anders Blok, and Sune Lehmann. Understanding scholar-trajectories across scientific
512 periodicals. *Scientific Reports*, 14(1):5309, 2024.

513

514 Computer Vision Foundation. Cvf open access. <https://openaccess.thecvf.com>, 2024.
515 Accessed: 2025-05-14.

516

517 Tirthankar Ghosal, Sandeep Kumar, Prabhat Kumar Bharti, and Asif Ekbal. Peer review analyze: A
518 novel benchmark resource for computational analysis of peer reviews. *Plos one*, 17(1):e0259238,
519 2022.

520

521 Miguel R Guevara, Dominik Hartmann, Manuel Aristarán, Marcelo Mendoza, and César A Hidalgo.
522 The research space: using career paths to predict the evolution of the research output of individuals,
523 institutions, and nations. *Scientometrics*, 109:1695–1709, 2016.

524

525 Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine Van Zuylen, Sebastian Kohlmeier,
526 Eduard Hovy, and Roy Schwartz. A dataset of peer reviews (peerread): Collection, insights and
527 nlp applications. *arXiv preprint arXiv:1804.09635*, 2018.

528

529 Jialiang Lin, Jiaxin Song, Zhangping Zhou, Yidong Chen, and Xiaodong Shi. Moprd: A multidis-
530 ciplinary open peer review dataset. *Neural Computing and Applications*, 35(34):24191–24206,
531 2023.

532

533 OpenReview. Openreview.net. <https://openreview.net>, 2024. Accessed: 2025-05-14.

534

535 Ana Carolina Ribeiro, Amanda Sizo, Henrique Lopes Cardoso, and Luís Paulo Reis. Acceptance
536 decision prediction in peer-review through sentiment analysis. In *EPIA Conference on Artificial
537 Intelligence*, pp. 766–777. Springer, 2021.

538

539 Google Scholar. Google scholar. <https://scholar.google.com/>. Accessed: 2025-05-14.

Nihar B Shah. Challenges, experiments, and computational solutions in peer review. *Communications
of the ACM*, 65(6):76–87, 2022.

David Soergel, Adam Saunders, and Andrew McCallum. Open scholarship and peer review: a time
for experimentation. 2013.

540 Shao-Hua Sun. Iclr2020-openreviewdata. [https://github.com/shaohua0116/](https://github.com/shaohua0116/ICLR2020-OpenReviewData/)
541 ICLR2020-OpenReviewData/, 2020. Accessed: 2025-05-14.

542

543 Jaroslaw Szumega, Lamine Bougueroua, Blerina Gkotse, Pierre Jouvelot, and Federico Ravotti. The
544 open review-based (orb) dataset: Towards automatic assessment of scientific papers and experiment
545 proposals in high-energy physics. *arXiv preprint arXiv:2312.04576*, 2023.

546 Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. Arnetminer: extraction and
547 mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD international
548 conference on Knowledge discovery and data mining*, pp. 990–998, 2008.

549

550 Claartje J Vinkenburg, Sara Connolly, Stefan Fuchs, Channah Herschberg, and Brigitte Schels.
551 Mapping career patterns in research: A sequence analysis of career histories of erc applicants.
552 *PloS one*, 15(7):e0236252, 2020.

553 Jing Yang. Position: The artificial intelligence and machine learning community should adopt a
554 more transparent and regulated peer review process. In *Forty-second International Conference on
555 Machine Learning (ICML) Position Paper Track*, 2025. URL [https://openreview.net/](https://openreview.net/forum?id=gnyqRarPzW)
556 forum?id=gnyqRarPzW.

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593