# Workshop on Reasoning and Planning for Large Language Models

### 1 Workshop Summary

The field of artificial intelligence is undergoing a profound transformation with the rise of large language models (LLMs), which are reshaping how we approach reasoning and planning. As these models grow in scale and complexity, their ability to not only process language but also apply it to reasoning-based tasks has emerged as a pivotal development. LLMs, such as OpenAI's o1 model, have shown remarkable provess in handling intricate tasks that require multi-step reasoning and long-horizon planning. These models excel at breaking down problems, generating detailed plans, and utilizing long-chain reasoning, marking a significant leap forward in AI's problem-solving abilities.

The integration of advanced reasoning processes within LLMs opens new opportunities for both research and practical applications. In response to these emerging needs, we propose the Workshop on Reasoning and Planning for LLMs. This workshop will explore the challenges of using Reinforcement Learning (RL) and other effective methods for training LLMs' reasoning and planning abilities, scaling inference for reasoning tasks, developing robust benchmarks for evaluation, and extending these models into multi-modal and embodied environments. As LLMs continue to expand their role in AI research and applications, this workshop will serve as a platform to discuss and refine the methods for advancing reasoning and planning within these models.

Overview of Discussion Topics. The workshop will cover a range of topics, including but not limited to:

(1) **Training Methodologies for Enhancing Reasoning and Planning Capabilities in LLMs:** We will explore the application of RL algorithms and other effective approaches in enhancing LLM reasoning and planning abilities during both pre-training and post-training stages. We will examine how techniques like Reinforcement Learning from Human Feedback (RLHF) can be adapted and expanded for efficient reasoning. Key questions include:

- How can RL and other effective methods be utilized in pre-training to improve reasoning abilities?
- What post-training approaches (e.g., fine-tuning, RLHF) are most effective for LLM planning tasks?
- How can synthetic data generation and self-supervised training enhance LLM reasoning and planning?

(2) Inference Time Scaling for Complex Reasoning Tasks: We will discuss challenges and innovations in scaling up reasoning during inference. As models become larger and tasks more complex, efficient inference mechanisms are critical. Topics of interest include

- What are the most promising methods for scaling inference times in reasoning-heavy tasks?
- How can models dynamically allocate resources during inference to optimize for reasoning and planning?

(3) Benchmarking Reasoning and Planning: Developing robust benchmarks for evaluating reasoning and planning in LLMs is critical to track progress. This session will address the need for new metrics and standardized tasks to assess reasoning abilities across different scenarios. Key discussions will include:

- What benchmarks can accurately reflect the reasoning and planning capabilities of LLMs?
- How do we design tasks that evaluate long-horizon reasoning and complex decision-making?

(4) **Multi-modality and Embodiment in LLMs:** As LLMs increasingly integrate with multi-modal environments, reasoning across multiple data types (e.g., vision, sound, text) becomes more essential. This session will explore the application of reasoning and planning in multi-modality and embodied AI systems, including robotics and real-world interactions:

- How can LLMs enhance multi-modal reasoning and planning to better interact with diverse environments?
- What are the key challenges and opportunities in applying LLMs to multi-modal tasks, including those requiring embodied reasoning?

(5) Exploring Broader Topics in Reasoning and Planning: In addition to the core themes mentioned above, our discussions will also encompass a broader range of emerging topics, including:

- Causal Reasoning: How can LLMs move beyond pattern recognition to infer causal relationships?
- Collaborative Reasoning in Multi-Agent Systems: How can LLMs enable multi-agent cooperation for distributed tasks?
- Uncertainty and Robustness: How can LLMs improve reasoning under ambiguous information?
- Human-in-the-Loop Systems: How can human feedback refine LLM decision-making processes?
- **Explainability:** How can we make LLM reasoning and planning more transparent and interpretable for real-world applications?

### 2 Invited Speaker Information

This workshop will feature senior scholars with extensive expertise in LLMs, RL, reasoning and decision-making, who will deliver keynote presentations about reasoning and planning of LLMs, and offer attendees valuable insights into the latest advancements in LLM reasoning and planning. The invited speakers are as follows:

**Sheila McIlraith, University of Toronto (Female, Confirmed).** Sheila McIlraith is a Professor in the Department of Computer Science, University of Toronto, a Canada CIFAR AI Chair (Vector Institute), and an Associate Director and Research Lead at the Schwartz Reisman Institute for Technology and Society. Prior to joining the University of Toronto, Prof. McIlraith spent six years as a Research Scientist at Stanford University, and one year at Xerox PARC. McIlraith's research is in the area of sequential decision making, broadly construed, with a focus on human-compatible AI. McIlraith is a Fellow of ACM, a Fellow of AAAI, Associate Editor for the Journal of AI Research (JAIR), and a past Associate Editor of the Journal of Artificial Intelligence. McIlraith is program Co-Chair of the 32nd AAAI Conference on Artificial Intelligence (AAAI), the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR2012), and the International Semantic Web Conference (ISWC2004). Her work on semantic web services has had notable impact. In 2011 she and her co-authors were honoured with the SWSA 10-year Award, recognizing the highest impact paper from the International Semantic Web Conference, 10 years prior.

**Guy Van den Broeck, University of California, Los Angeles (Male, Confirmed).** Guy Van den Broeck is a Professor and Samueli Fellow at UCLA, in the Computer Science Department, where he directs the Statistical and Relational Artificial Intelligence (StarAI) lab. His research interests are in Machine Learning, Knowledge Representation and Reasoning, and Artificial Intelligence in general. His papers have been recognized with awards from key conferences such as AAAI, UAI, KR, OOPSLA, and ILP. Guy is the recipient of an NSF CAREER award, a Sloan Fellowship, and the IJCAI-19 Computers and Thought Award.

**Hongyu Ren, OpenAI (Male, Confirmed).** Hongyu is a research scientist at OpenAI exploring ChatGPT. He is the core contributor to OpenAI o1 and the co-leader of GPT-40 mini and OpenAI o1 mini. He received his Ph.D. in Computer Science from Stanford University and a B.S. with honors in Computer Science from Peking University. He is also a recipient of the Apple PhD Fellowship, Baidu PhD Fellowship, and Masason Foundation Fellowship.

**Timothy P. Lillicrap, Google DeepMind & University College London (Male, In Contact).** Timothy is a staff research scientist at Google DeepMind, and adjunct professor at University College London. He has contributed to AlphaGo and AlphaZero, specializing in machine learning, reinforcement learning, and decision-making frameworks. His work bridges AI and neuroscience, with notable contributions to deep neural networks and one-shot learning. Lillicrap holds a B.Sc. in cognitive science and AI from the University of Toronto and a Ph.D. in systems neuroscience from Queen's University. He joined DeepMind in 2014 and became a staff research scientist in 2016. He has received multiple honors, including an NSERC Fellowship and European Research Council grants.

**Yuandong Tian, Meta (Male, Confirmed).** Yuandong Tian is a Research Scientist Director in Meta AI Research (FAIR), leading the group of reasoning, planning and decision-making with Large Language Models (LLMs). He is the project lead for OpenGo project that beats professional players with a single GPU during inference, serves as the main mentor of StreamingLLM and GaLore that improve the training and inference of LLM, and is the first-author recipient of 2021 ICML Outstanding Paper Honorable Mentions DirectPred and 2013 ICCV Marr Prize Honorable Mentions HierarchicalDataDrivenDescent, and also received the 2022 CGO Distinguished Paper Award CompilerGym. Prior to that, he worked in Google Self-driving Car team in 2013-2014 and received a Ph.D in Robotics Institute, Carnegie Mellon University in 2013.

Natasha Jaques, University of Washington & Google DeepMind (Female, Confirmed). Natasha is an Assistant Professor at the University of Washington's Paul G. Allen School of Computer Science & Engineering and a Senior Research Scientist at Google DeepMind, where she leads the Social Reinforcement Learning Lab, focusing on integrating social learning and multi-agent training to enhance AI agents' learning, coordination, and human-AI interaction. She is also a Senior Research Scientist at Google DeepMind. Her PhD work at MIT contributed to fine-tuning language models with RL and learning from human feedback, which influenced OpenAI's RLHF. She has interned at DeepMind, Google Brain, and served as an OpenAI Scholars Mentor. She was a Visiting Postdoctoral Scholar at UC Berkeley and a Senior Research Scientist at Google Brain, where she advanced methods for adversarial environment generation to improve RL agent robustness. Her work has earned several awards, including Best Demo at NeurIPS and an ICML Best Paper Honourable Mention.

**Bo An, Nanyang Technological University (Male, Confirmed).** Bo is a President's Chair Professor and Head of Division of Artificial Intelligence at the College of Computing and Data Science of the Nanyang Technological University (NTU). He is also Co-Director of Artificial Intelligence Research Institute (AI.R) of NTU. During October 2010 to June 2012, he was a Postdoctoral Researcher at the University of Southern California, working with Professor Milind Tambe. He received the Ph.D degree in Computer Science from the University of Massachusetts, Amherst, where he was advised by Professor Victor Lesser. His research focuses on multi-agent systems, game theory, reinforcement learning, negotiation, and optimization. He was named to IEEE's "AI's 10 to Watch" list in 2018. He serves on several editorial boards and was PC Co-Chair of AAMAS'20 and General Co-Chair of AAMAS'23. He was invited to give a keynote speach at IJCAI 2024 and will be PC Chair of IJCAI'27

**Stephen McAleer, OpenAI (Male, Confirmed).** Stephen is a research scientist at OpenAI for AI agent safety and reasoning. Previously, he was a postdoc at CMU working with Tuomas Sandholm. He received his PhD in computer science from the University of California, Irvine working with Pierre Baldi. During his PhD, he did research scientist internships at Intel Labs and DeepMind.

**Junxian He, Hong Kong University of Science and Technology (Male, Confirmed).** Junxian is a tenure-track assistant professor in The Hong Kong University of Science and Technology, Department of Computer Science and Engineering. He obtained his PhD from Carnegie Mellon University, Language Technologies Institute in 2022, where he was co-advised by Graham Neubig and Taylor Berg-Kirkpatrick. Before that, he received the bachelor degree in Electronic Engineering from Shanghai Jiao Tong University in 2017. he also spent some time at Facebook AI Research (2019) and Salesforce Research (2020).

### **3** Tentative Schedule and Organizational Details

**Tentative Schedule.** Our workshop will be conducted in a hybrid format, accommodating both in-person and virtual participation. The program will feature six invited talks, each lasting 30 minutes (25 minutes for the presentation and 5 minutes for Q&A), along with five contributed talks of 15 minutes each, selected from submitted papers. Additionally, there will be two 1-hour poster sessions and a 1-hour panel discussion focused on the future of reasoning and planning for LLMs. The poster sessions will be accessible both virtually and in person. Virtual participants can select a session based on their time zone, with dedicated virtual spaces hosting the poster presentations. To foster engagement and interaction among remote attendees, a dedicated chat platform will be available, ensuring seamless communication throughout the workshop.

#### Morning:

#### Afternoon:

- 08:50 09:00 Introduction and opening remarks
- 09:00 09:30 Invited Talk 1
- 09:30 10:00 Invited Talk 2
- 10:00 11:00 Poster Session 1
- 11:00 11:30 Invited Talk 3
- 11:30 12:00 Invited Talk 4
- 12:00 13:30 Break

- 13:30 14:00 Invited Talk 5
- 14:00 14:45 Contributed Talk 1 & 2 & 3
- 14:45 15:45 Poster Session 2
- 15:45 16:15 Invited Talk 6
- 16:15 16:45 Contributed Talk 4 & 5
- 16:45 17:45 Panel discussion
- 17:45 18:00 Paper Award and Closing Remarks

Important Dates for Review Process. We will follow the suggested dates by ICLR.

- Workshop paper submission deadline: February 3, 2025.
- Workshop paper notification date: March 3, 2025
- Final workshop program, camera-ready, videos uploaded: April 3, 2025

**Support for Tiny or Short Papers.** In addition to full-length papers (6-8 pages), we will follow the ICLR workshop guidelines by offering a Tiny/Short Papers track. This track will include 2-4 page submissions focused on implementations and evaluations of unpublished but simple ideas, moderate yet self-contained theoretical results, follow-up experiments or re-analyses of published work, or fresh perspectives on existing publications. This format provides an opportunity for underrepresented, resource-constrained, or early-career researchers to receive feedback and engage in discussions at ICLR.

Anticipated Audience Size. Our workshop is designed for researchers and practitioners with a strong interest in machine learning and natural language processing, with a focus on LLMs reasoning and planning. Based on the success of previous workshops on related topics, the relevance of the proposed agenda, and the appeal of our invited speakers, we anticipate a consistent in-person attendance of 400-500 participants, with a total reach of 600-700 attendees over the course of the event.

Access, Broadcast, and Promotion. We will maintain a dedicated website for the workshop and actively promote it to a broad audience. For invited talks and panel discussions, we will post the abstracts and speaker biographies as soon as they are confirmed, which is expected to be several months ahead of the event. After the workshop, we will strive to make slides and recordings of the talks available online. The workshop will be conducted in a hybrid format and live-streamed throughout the day. Both speakers and participants will have the option to join either in person or online, with provisions to ensure virtual attendees can ask questions and participate in discussions. Accepted papers will be posted on the workshop website, with an option to publish them on arXiv, noting their association with the event. We will ensure accessibility for virtual attendees by providing captions and other necessary accommodations. For promotion, in addition to the website, we will leverage social media platforms including X and LinkedIn to reach a wider audience. Organizers and speakers will be encouraged to share the event within their networks. We will also use mailing lists, such as those available through OpenReview, to reach relevant researchers in reasoning and planning, sending group emails to introduce the workshop and invite participation.

**Sponsorship.** We have engaged in discussions with Microsoft, Google, Salesforce, ByteDance, and NUS, all of whom have expressed interest in providing sponsorship support. The specific level of support will be finalized once the workshop is approved. Potential sponsorship will include funding for speaker honorariums, support for underrepresented groups, and coverage of registration fees for some students, early-career researchers, and participants from low-income regions. Additionally, sponsorship will assist in covering F&B costs during the event.

## 4 Differences with Previous Workshops

Several recent workshops have explored themes relevant to our focus, including the "LLM Agents Workshop" (ICLR 2024), "System 2 Reasoning At Scale" (NeurIPS 2024), "Knowledge and Logical Reasoning in the Era of Data-driven Learning" (ICML 2024), "Natural Language Reasoning and Structured Explanations Workshop" (ACL 2024), "3rd Workshop on the Interactions between Analogical Reasoning and Machine Learning" (IJCAI 2024), and "Neural Reasoning and Mathematical Discovery" (AAAI 2025).

Our workshop distinguishes itself by addressing the practical challenges of reasoning and decision-making in complex environments, with a particular focus on reinforcement learning techniques and other methods to enhance the reasoning and planning capabilities of LLMs. While other workshops take a broader view of reasoning, we concentrate on integrating reinforcement learning for long-horizon planning and post-training optimization. Additionally, we will develop robust benchmarks specifically designed to evaluate LLM reasoning and planning performance at scale, offering a more practical, application-driven approach. Lastly, our workshop prioritizes extending LLMs into multi-modal and embodied environments, enabling more advanced reasoning and planning in real-world, physical contexts.

### **5** Diversity Statement

**Diversity of the Organizing Committee.** Our organizing committee is composed of individuals from diverse demographic backgrounds and areas of expertise. We are committed to fostering diversity across several dimensions, including gender, seniority, professional experience, institutional affiliation, and research focus. Of our 9 organizers, 2 identify as female and 7 as male. The committee includes 3 assistant professors, 2 associate professors, 2 industry research scientists, and 2 senior PhD candidates, representing institutions across the globe including Google Deep-Mind, Yale University, University of Washington, the Allen Institute for AI, National University of Singapore, and King's College London. Each organizer holds rich experience as a leader or invited speaker in prominent AI and deep learning workshops.

**Diversity of Speakers.** We are equally committed to diversity in the selection of invited speakers, who represent leading institutions worldwide, such as Google DeepMind, OpenAI, Meta, University of Toronto, University of Washington, University of California, Los Angeles, Nanyang Technological University, and Hong Kong University of Science and Technology. Our speakers range from full professors to industry leaders, offering a broad mix of academic and industry insights. Notable speakers include Sheila McIlraith, a fellow of the ACM and AAAI, Guy Van den Broeck, a Sloan Fellowship recipient, and Hongyu Ren, a core contributor to OpenAI's o1 project and co-leader of GPT-40 mini. We are committed to gender representation and ensuring diverse perspectives are highlighted. As we finalize our speaker lineup, we remain focused on achieving balanced representation across the global AI community.

**Diversity of Participants.** Our workshop aims to attract a diverse group of participants, reflecting the wide-ranging interest in reasoning and planning for LLMs. We are actively encouraging participation from individuals with varied backgrounds and disciplines. To support underrepresented groups, we will offer sponsorship to facilitate attendance. In addition, we will provide registration fee support for students, early-career researchers, and participants from low-income regions, along with the option for virtual attendance. These efforts are designed to promote accessibility and ensure that diverse perspectives contribute to the workshop.

By fostering an inclusive and welcoming atmosphere, we aim for our diverse organizing committee and invited speakers to inspire attendees from all backgrounds, encouraging vibrant and productive discussions.

### 6 Organizer Information

**Pang Wei Koh** (University of Washington, Male). Pang Wei is an assistant professor in the Allen School of Computer Science and Engineering at the University of Washington, a visiting research scientist at AI2, and a Singapore AI Visiting Professor. His research interests are in the theory and practice of building reliable machine learning systems. His research has been published in Nature and Cell, featured in media outlets such as The New York Times and The Washington Post, and recognized by the MIT Technology Review Innovators Under 35 Asia Pacific award and best paper awards at ICML and KDD. He received his PhD and BS in Computer Science from Stanford University. Prior to his PhD, he was the 3rd employee and Director of Partnerships at Coursera. *He organized the Workshop on Distribution Shifts at NeurIPS 2021–2023, Workshop on Navigating and Addressing Data Problems for Foundation Models and Workshop on Reliable and Responsible Foundation Models at ICLR 2024*.

**Arman Cohan (Yale University, Male).** Arman is an Assistant Professor of Computer Science at Yale University and a faculty Research Scientist at the Allen Institute for AI (Ai2). His research spans various problems at the intersection of Machine Learning and Natural Language Processing, including Language Modeling, Representation Learning, and applications to specialized domains including science. Prior to joining Yale, he was a Research Scientist at AI2 and an Affiliate Assistant Professor at University of Washington. His research has been recognized with multiple awards including a best paper at EMNLP and ACL, outstanding paper at EACL, and an honorable mention at COLING. *He is the Publications Chair for NAACL 2025 and previously organized the Workshop on SciNLP - Scientific NLP at AKBC 2020, and Workshop on Scholarly Document Processing at NAACL 2021 and COLING 2022.* 

**Xidong Feng (Google DeepMind, Male).** Xidong is a research scientist at Google DeepMind. His research spans over Large Language Model, Reinforcement Learning and Multi-agent Learning. He has published over 10 papers in top AI conferences or journals like NeurIPS, ICML and JMLR. He obtained his Ph.D. at Computer Science, University College London, advised by Prof. Jun Wang. He previously earned his BS degree from Tsinghua University.

**Min-Yen Kan (National University of Singapore, Male).** Min-Yen is an Associate Professor with research interests spanning digital libraries, natural language processing, and language reasoning. A senior member of both the ACM and IEEE, he completed his Ph.D. at Columbia University. Additionally, he holds the role of Vice Dean at the NUS School of Computing. *He currently serves as the ACL Ethics Co-Chair and was General Co-Chair for COLING 2024.* 

**Bryan Hooi (National University of Singapore, Male).** Bryan is an Assistant Professor in the School of Computing and the Institute of Data Science in National University of Singapore. His research aims to make machine learning systems more reliable and applicable to a wider variety of real-world contexts, particularly. He received his Ph.D. in Machine Learning from Carnegie Mellon University, advised by Christos Faloutsos. He received his M.S. in Computer Science and his B.S. in Mathematics from Stanford University. His research is recognized as best paper awards at KDD 2016 and ICDM 2020. *He was the organizer of The WebConf 2024 Workshop on Graph Foundation Models* 

**Nouha Dziri (Allen Institute for AI, Female).** Nouha is a research scientist at the Allen Institute for AI (AI2), working with Yejin Choi. She co-leads efforts to build OLMo, an open LLM focused on advancing AI. Her research centers on understanding the limits of Transformers, their inner workings, and alignment in LLMs. Her work has been featured in publications such as TechCrunch, LeMonde, The Economist, and Science News. *She is the Demo Track Chairs for NAACL 2025.* 

Yali Du (King's College London, Female). Yali Du is a Senior Lecturer (Associate Professor) in AI at King's College London, and a Turing Fellow at The Alan Turing Institute. She leads the Cooperative AI Lab. Her research aims to enable machines to exhibit cooperative and safe behaviour in intelligent decision making tasks. She was chosen for the AAAI New Faculty Highlights award (2023), Rising Star in AI 2023. *She also serves in organising committee for AAMAS 2023 (Local Arrangement Team) and NeurIPS 2024 (Communication Chair)*.

**Zhiyuan Hu** (National University of Singapore, Male). Zhiyuan is a PhD candidate at the National University of Singapore, specializes in LLM reasoning and planning. He has published over 10 papers in top AI conferences like AAAI, ACL, ICLR and NeurIPS, and served on program committees for NLP and ML conferences. His past roles include a visiting researcher at University College London's AI Centre. He is a recipient of the NUS President's Graduate Fellowship. *He also serves the chair in NLP SIG, NUS*.

**Yilun Zhao (Yale University, Male).** Yilun is a PhD student in the Yale NLP group, advised by Professor Arman Cohan. His research focuses on multimodal reasoning and enhancing LLMs for scientific workflow. He has published over 20 papers in top AI conferences like ACL, NeurIPS, EMNLP, and NAACL. *He serves the student chair in SciNLP 2024, Yale.* 

Simeng Han PhD Student, Yale University	Gabrielle Kaili-May Liu PhD Student, Yale University	Jacob Dunefsky PhD Student, Yale University	Yitao Long PhD Student, New York University
Hongjun Liu PhD Student, New York University	Daniel Shao Master Student, Yale University	<b>Evan Shi</b> Master Student, Yale University	<b>Guo Gan</b> Master Student, Zhejiang University
<b>Chunyuan Deng</b> PhD Student, Rice University	<b>Roy Xie</b> PhD Student, Duke University	Manasi Patwardhan Research Scientist, TCS Research	Aniket Deshpande Research Scientist, TCS Research
Zhenting Qi Master Student, Harvard University	Lin Xu PhD Student, National University of Sin- gapore	<b>Jiaying Wu</b> Postdoc Fellow, National University of Singapore	Weihao Tan PhD Student, Nanyang Technological University
Longtao Zheng PhD Student, Nanyang Technological University	Iman Mirzadeh ML Research Engineer, Apple	<b>Oliver Slumbers</b> PhD student, University College London	Haotian Fu PhD student, Brown University
Mengyue Yang Assistant Professor, University of Bristol	<b>Jean-Francois Ton</b> Senior Research Scientist, ByteDance Research	Muhammad Faaiz Taufiq PhD Student, University of Oxford	Zhenghai Xue PhD Student, Nanyang Technological University
Wang Suyuchen PhD Student, Mila	Jakub Grudzien Kuba PhD Student, UC Berkeley	Benjamin Liu PhD Student, National University of Sin- gapore	Shuofei Qiao PhD Student, Zhejiang university
Yong Jiang Research Scientist, Alibaba Group	Xiang Chen PhD Student, Zhejiang university	Ningyu Zhang Associate Professor, Zhejiang Univer- sity	<b>Thong T. Nguyen</b> PhD Student, National University of Sin- gapore
Anh Tuan Luu Assistant Professor, Nanyang Techno- logical University	See Kiong Ng Professor, National University of Singa- pore	Yubin Kim PhD student, Massachusetts Institute of Technology	Bryan Hooi Assistant Professor, National University of Singapore
Anirudh Goyal Research Scientist, Google DeepMind	Liangming Pan Assistant Professor, University of Ari- zona	Sachin Goyal PhD Student, Carnegie Mellon Univer- sity	Yue Feng Assistant Professor, University of Birm- ingham

Table 1: Pool of Candidate Reviewers

### 7 Program Committee and Reviewing

As detailed in Table 1, we have invited 40 researchers from 23 institutions worldwide to serve on our Program Committee. This committee comprises a diverse and well-rounded group of ML/NLP experts at various career stages, including graduate students, postdoctoral researchers, and faculty members. Each member has at least three publications in top-tier ML/NLP conferences and has already confirmed their participation as a reviewer.

Based on our previous experience in workshop organization, we anticipate receiving 60-80 paper submissions. The initial review of these submissions will be handled by the Program Committee through a double-blind process. We will continue to invite additional reviewers as needed to ensure that each paper receives at least three reviews, with a maximum of three papers assigned to any individual reviewer.