

CLAI Unconf - Call for Pre-registered Papers

The natural world is a *dynamic* environment that is constantly changing, requiring humans and animals to *adapt*. However, much research in Artificial Intelligence (AI) is focused on the static train-test paradigm with heavily curated benchmarks. To transcend these static assumptions, the Continual Artificial Intelligence Unconference 2023 (CLAI Unconf) invites submissions with original contributions in the form of *pre-registration articles* on topics broadly related to continual learning and adjacent fields. Solicited themes may revolve around (but are not limited to) navigating complex data collection systems, understanding and describing continuous streams of data, lifelong learning processes and generalization of knowledge beyond a specific target, discovering new concepts in changing environments, and related topics from an interdisciplinary perspective.

Pre-registration ensures scientific excellence through a two-stage submission procedure that separates quality of scientific hypotheses from potentially negative downstream empirical analysis. In the first stage, researchers write a proposal with well-articulated ideas and a thoroughly outlined experimental protocol, which is peer-reviewed on the OpenReview platform. Upon acceptance, researchers can incorporate feedback through active dialogue at the conference. In the second stage, researchers will rigorously evaluate their ideas, with the final paper published in CLAI Unconf's Proceedings of Machine Learning Research ([PMLR](#)), after a light second round of peer-review to confirm rigorous execution. In this way, ideas can be presented and discussed within the community before being validated experimentally, allowing valuable insights from well-formulated hypotheses, even if later they do not empirically outperform other methods. This should help move the field towards papers with high-quality hypotheses and away from an increasingly positive-result bias in published literature.

We welcome papers across many disciplines, including but not limited to more specific topics within:

- (Deep) learning beyond static datasets: deep continual representation learning and foundation models, distributed learning over time and multi-agent systems, semi-supervised and unsupervised continual learning, meta-learning and continual reinforcement learning;
- Online learning and learning in non-stationary environments: adaptation and knowledge accumulation, learning with fixed features, computational and memory efficiency;
- Discovery of novel elements over time: out-of-distribution detection, drift detection, open world learning, handling partially observable information from potentially disparate sources;
- Continual learning in the brain: neuroscientific perspectives, memory consolidation and learning in the brain, new insights from and for cognitive science, bio-inspired methods;
- Practical continual learning applications: robotics, vision, language learning, predictive maintenance, AutoML and MLOps, continuum edge-cloud; and
- Critical analyses & positions: evaluation, experimental protocols, fairness & ethical considerations, terminology consolidation, documentation, complex data collection systems.

We emphasize that the pre-registration process moves the priority of reviewing to gauging *technical correctness, quality of the scientific hypotheses and validity of the outlined experimental protocol*. Upon careful review of the soundness of the pre-registered proposal, accepted papers are virtually guaranteed to disseminate their findings (no matter their nature), as long as they *rigorously follow their protocol and adequately detail their insights*. The format is thus particularly suitable for highly novel and risky ideas: there is less focus on beating state-of-the-art, and less worrying about whether sensible conjectures will yield negative results. We equally solicit submission of critical

analyses that highlight their necessity and outline how the ensuing study will increase understanding during pre-registration.

Important Dates

First stage

July 21: pre-registered paper submission
August 20: reviews released
August 30: author rebuttals due
September 7: decisions communicated
September 20: camera ready due
October 19/20: CLAI unconference

Second stage

April 16, 2024: submission of final papers
May 18, 2024: decisions released
June 5, 2024: journal camera ready due

Pre-registration and Review Process

Across an increasing number of scientific fields, pre-registrations and registered reports have become commonplace to improve the rigor and incentives of research publications. Our pre-registration process is similar to a previous NeurIPS workshop, see <https://preregister.science/> for paper examples. For a more general understanding across fields, see <https://www.cos.io/initiatives/registered-reports>.

Pre-registered proposals provide all methods and hypotheses for the proposed study for peer-review and acceptance, prior to when experiments are conducted. The emphasis of peer-review is thus placed upon the ideas put forward, and whether the proposed experiments are adequate to meet the studies' aims. This review method thus helps both the formulation of clear and plausible research conjectures, as well as to mitigate the positive-results bias in the literature. *Respectively, both negative and positive results are published alike.* The rigorously outlined experimental protocol further promotes reproducibility, which is presently a general concern across all sciences and AI research.

The format is particularly suitable to continual learning, as pre-registration also ensures authors run experiments that the community sees value in. A widely-accepted problem in continual learning is that papers often make ad-hoc assumptions in their experimental setup, making it difficult to compare between methods. Although differences are expected in a recent and fast-growing field such as continual learning, pre-registration should encourage experiments to be more standardized, or at the very least, ensure they are compared transparently and fairly.

While registered reports add substantial gains in scientific rigor, they add minimal additional overhead (despite what some might initially think), effectively splitting the research process into two stages:

First stage:

- 1.) In the first stage, a proposal paper is written for review *outlining all background, proposed methods, hypotheses and a carefully thought out experimental protocol.*

- 2.) This proposal is then thoroughly peer-reviewed in a double-blind process, judging the technical correctness, quality of presentation, soundness of the hypothesis, rigor and feasibility of the outlined experimental protocol.
- 3.) Upon acceptance, the proposal's ideas are presented at CLAI Unconf to share hypotheses and planned experimentation, as well as actively gather feedback.

Second stage:

- 4.) In a second stage, the proposed research is conducted and the results are documented for publication in an additional experiments section of the paper. The earlier pre-registered parts of the paper remain untouched, with only minor edits being allowed to summarize the final findings in the abstract/introduction.
- 5.) A lighter final single-blind peer review (anonymity of reviewers is preserved) then accepts this full paper version, on the basis of whether the protocol has been followed rigorously and findings described appropriately. Whether these findings are confirmatory of the hypotheses does not affect the outcome.
- 6.) The final complete paper is published in PMLR.

In summary, authors *first finalize all methods, hypotheses and experimental protocols, which are double-blind peer-reviewed. Upon acceptance, authors will be allotted time to execute their outlined studies. A second, lighter, peer-review step confirms that the proposal has rigorously been adhered to.* We will run the full submission process on OpenReview and ensure the proposal reviewing process is double-blind, with the second stage review then being single-blind for the final executed result papers.

Submission Formatting

Pre-registered proposals are expected to be 6 pages in length, outlining all background, proposed methods, hypotheses and a carefully thought out experimental protocol, with an unlimited number of pages for references. Whereas we don't strictly impose a structure, the above elements need to be clearly outlined and we strongly encourage authors to adopt the following sections: introduction, background and related work, hypotheses and proposed method, experimental protocol.

Upon acceptance of the pre-registered proposal, the final paper is required to add an experimental results and discussion section, *with no strict page limit.* Importantly, the *initial six pages are required to be left unmodified*, with the exception of being allowed a minor addition of a summary of results at the end of the abstract and/or introduction. The unlimited amount of pages is to ensure that all results are properly documented and adequately explained. However, we encourage authors to stay concise, with an envisioned final length of 10-12 pages. Additional materials (e.g. videos, code) are permitted in the form of supplementary material in this stage. We strongly encourage the contribution of public code in the second stage to promote reproducibility.

Papers need to be formatted according to the CLAI Unconf LaTeX template, derived from PMLR, which will be provided soon. Any modifications to the style file, in particular the alteration of margins, are not permitted and will result in desk rejection. Similarly, dual-submissions are not allowed, and papers will face desk rejection if they are concurrently being submitted to another venue. Submissions based on prior non-arXiv works (e.g. workshops without proceedings) are

allowed, but we emphasize that they need to comply with the spirit of pre-registration and should not already have empirical evidence to fully confirm the main hypotheses.

Anonymization and Author Order

The initial submission for the pre-registration is required to be fully anonymized. Any references to the authors' prior work need to adhere to the double-blind policy and are not allowed to expose the authors' identities or affiliations. Upon acceptance, the later final reports are submitted for single-blind review, where the authors' identities are known (given that they have presented at the conference).

Importantly, we do not permit the addition or removal of authors after the initial pre-registered submission's deadline (reordering of authors from proposal to final stage papers may be allowed in exceptional circumstances upon written request to the chairs). As pre-registration aims at promoting scientific rigor, this policy reflects the authors' commitment to sound scientific practice.

Second Stage Protocol Deviations

Minor changes to the experimental protocol in the second stage will be permitted, e.g. in the case that an experiment has been added to provide additional insights that have been explored during experimentation, as long as the main protocol has been rigorously exercised. In this case, a mandatory final section called "Changes to the protocol" is required. We strongly urge authors to carefully consider whether any alterations to the experimental protocol are in fact necessary and highly recommend changes be limited exclusively to additional content. Outright removal of essential initially-proposed parts is likely to result in a second stage rejection, as the format explicitly encourages the reporting of negative results.

Attendance

CLAI Unconf is a fully virtual event with multiple time zones to accommodate participants from all around the world. Authors whose pre-registered proposals have been accepted agree to present their works as either a contributed talk (if selected) or in a poster session virtually (GatherTown).