Continual Causality: A Retrospective of the Inaugural AAAI-23 Bridge Program

Martin Mundt *TU Darmstadt, hessian.AI, ContinualAI* MARTIN.MUNDT@TU-DARMSTADT.DE
Keiland W. Cooper *University of California, ContinualAI* KWCOOPER@UCI.EDU
Devendra Singh Dhami *TU Darmstadt, hessian.AI* DEVENDRA.DHAMI@TU-DARMSTADT.DE
Adèle Ribeiro *University of Marburg* ADELE.RIBEIRO@UNI-MARBURG.DE
James Seale Smith *Georgia Tech, ContinualAI* JAMESSEALESMITH@GATECH.EDU
Alexis Bellot *DeepMind* ABELLOT@DEEPMIND.COM
Tyler Hayes[*] *Rochester Institute of Technology, ContinualAI* TLH6792@RIT.EDU

Abstract

Both of the fields of continual learning and causality investigate complementary aspects of human cognition and are fundamental components of artificial intelligence if it is to reason and generalize in complex environments. Despite the burgeoning interest in investigating the intersection of the two fields, it is currently unclear how causal models may describe continuous streams of data and vice versa, how continual learning may exploit learned causal structure. We proposed to bridge this gap through the inaugural AAAI-23 "Continual Causality" bridge program, where our aim was to take the initial steps towards a unified treatment of these fields by providing a space for learning, discussions, and to build a diverse community to connect researchers. The activities ranged from traditional tutorials and software labs, invited vision talks, and contributed talks based on submitted position papers, as well as a panel and breakout discussions. Whereas materials are publicly disseminated as a foundation for the community: https://www.continualcausality.org, respectively discussed ideas, challenges, and prospects beyond the inaugural bridge are summarized in this retrospective paper.

1. Continual Causality: Bridging Continual Learning with Causality

For an agent to successfully learn in natural settings, it must accommodate often non-linear, emergent dynamics from a sea of apparent stochasticity. For machine learning agents, the last decade has seen spectacular progress in these settings, surpassing human performance across a host of complex tasks. However, it is not only essential to these problems to identify the dependence between variables in observational data and their dynamics, but it is also essential to understand their underlying causation and how their dependencies arise.

Causality theory (Pearl, 2009; Peters et al., 2017) provides language, algorithms, and tools to discover and infer cause-and-effect relationships from any collection of observational/experimental data based on a partial understanding of a complex system. The field provides a powerful framework for learning a (partial) representation of the underlying causal model, evaluating the effect of unrealized interventions, reasoning about the effects of stochastic policies, generalizing causal knowledge to a target population, recovering from

[*] Now at NAVER LABS Europe

selection bias, and deriving counterfactual explanations. Moreover, there is increasing interest in learning and leveraging causal knowledge in dynamic environments with shifts in data distribution (Tian, 2008; Bareinboim and Pearl, 2016; Davis et al., 2018).

In machine learning, many of these problems appear under the rubric of continual learning. Continual learning systems (Chen and Liu, 2018; Hadsell et al., 2020) learn over time from a continuous stream of data, enable knowledge transfer and alleviate potentially malicious interference when distributional shifts are experienced. At the crux, they achieve the latter without revisiting previously encountered training samples (unless they are in memory), much in contrast to traditional machine learning settings where all training samples are available from the start and can be revisited.

The fields of causality and continual learning thus exhibit complementary strengths. On one hand, causal models formally describe structural similarities between domains that can aid the generalization of continual learning systems. On the other hand, continual learning systems efficiently navigate complex data collection and model streams of data that have not yet been studied from a causal perspective. These are both unavoidable components to move towards the goal of building autonomous or 'human-like' systems (Lake et al., 2017). The latter should be capable of learning continuously in the presence of different types of data and performing cause-and-effect reasoning about its surrounding world.

1.1. Goals: challenges and opportunities

Despite causality having recently taken huge strides, causal learning tools are prone to catastrophic interference of past-learned and possibly imperfect representations when new data becomes available. Learning (possibly dynamic and out-of-equilibrium) causal models from a continuous stream of data is thus extremely challenging, especially when data comes from multiple disparate sources and is collected under different conditions and study designs.

The field of continual learning is poised to mitigate these shortcomings. Despite prior efforts to transfer causal knowledge across domains and populations (Pearl and Bareinboim, 2014; Magliacane et al., 2018; Lee et al., 2020; Teshima et al., 2020), a challenging assumption of sufficient background knowledge on what features are shared across the domains is made. To aid causal learning tools, the addition of continual learning may mitigate the fragility of causal models to data distribution shifts with partially available information, as well as provide benefits to learning from temporal structure of streams of data.

In turn, causal tools present an opportunity to aid continual learning algorithms. Despite continual learning research considering a plethora of practical measures (e.g. memory, compute, model size) and nuanced temporal assessment (e.g. forward transfer, backward transfer, forgetting), the causal structure seems to be yet missing from the picture (see e.g. the evaluation and desiderata of Kudithipudi et al. (2022); Mundt et al. (2022)). Yet, causal models can provide a framework to improve a continual learner's capabilities by discovering and describing the underlying system beyond the statistical relationships learned from a stream of data received from the environment. Thus, causality presents an opportunity to further improve not only a continual model's explainability and effectiveness in decision making, but also, by placing causal modeling in a more naturalistic, continuous setting, our human understanding of continual learning and causality alike.

Despite a few recent attempts (Javed et al., 2020; Chu et al., 2020; Gong et al., 2022; Kancheti et al., 2022), it is presently not clear how to incorporate the advances of causal inference with those of continual learning and vice versa. Thus, the interface between these fields remains a relatively unexplored research area. It is our thesis of the bridge program that by placing these disciplines under a conceptual and theoretical umbrella, their intersection will help catalyze advances in prediction and explanation in machine learning, thus taking another step towards building general-purpose artificial intelligence. Through the bridge program, the first steps towards a unified treatment have been taken.

1.2. Specific initial open questions

Causal modeling and inference are well aligned with human-like cognition in terms of i) association ("recognition"), ii) intervention ("thought experiments"), and iii) counterfactuals ("reflection"). Due to this alignment, causal models are vulnerable to pitfalls that humans suffer from, such as limited memory, leading to difficulties in learning multiple concepts continuously. Since continual learning systems learn from continuous streams of data, answering the question of *"how to continually learn and exploit causal systems in dynamic non-stationary environments"* has been posed as the initial major focus of the inaugural bridge. The question naturally allows for a chain of interesting thoughts and immediate outcomes, with the prospect of mutually advancing the fields. The following imminent research items have been formulated, comprised of aspects that were expected to start the continual causality bridge, opening up the potential for more advanced questions later:

1. Define and react to catastrophic interference and knowledge transfer in learning causal models in the context of a continuous, non-stationary stream of data.

2. Understand effective ways for the causal structure to aid in leveraging the accumulated knowledge of a continual learning system and interpret distributional shifts.

3. Develop next-generation benchmarks that go beyond the re-purposing of existing datasets to adequately support the above items and further essential research questions towards a symbiosis of continual learning and causality.

2. The Continual Causality Bridge

As a two-day event, the Continual Causality bridge offered a varied set of activities that have supported education and discussion at various levels of expertise. Educational activities involved two traditional tutorials, as an entry to the respective fields of continual learning and causality, for both newcomers and experts that are only familiar with either of them. Similarly, two software labs have provided respective hands-on experience, based on the popular emerging software tools *Avalanche*[1] (Lomonaco et al., 2021) and *DoWhy*[2].

To complement these educational activities, a set of vision and contributed talks have provided initial positions on how the fields of continual learning and causality can be brought further together, why a necessity for the latter exists, and what respective challenges may

1. https://avalanche.continualai.org
2. https://py-why.github.io/dowhy

need to be faced in the imminent future. Finally, the invited vision talks and contributed position works have served as a basis for a panel and interactive breakout sessions. Participants of the bridge have not only received the chance to network but also to discuss presented ideas and to delve into a deeper exchange in a more detailed conversation. The various outcomes of these sessions have been collected and are disseminated in the final section of this retrospective paper, serving the role of a foreword to the bridge's full proceedings, see next section.

2.1. Call for papers

Following the bridge's spirit of presenting visions and tangible ideas to connect the fields of continual learning and causality, a call for papers had been issued several months prior to the inaugural event. Two-page paper contributions (AAAI dual-column format, excluding references) were solicited to outline challenges that need to be overcome, highlight synergies, and propose future steps. Following a rigorous double-blind peer-review review process on OpenReview, authors were allowed to incorporate any received feedback for the proceedings in an extra page, resulting in up to six pages in the single-column proceedings format.

The rationale behind the call for papers was for prospective community members to voice diverse views that have the potential to advance AI through an ongoing cross-disciplinary exchange. We have therefore not imposed strict constraints on the exact sub-topics of submissions, as long as they had targeted the overall goal of bridging the fields. Submissions were free to focus on a single particularly interesting synergy, take an angle from a specific scientific discipline, or sketch a grander view. To provide some initial inspiration, a few pointers had been given, including but not limited to:

- Continual learning and exploitation of causal systems in dynamic non-stationary environments
- Catastrophic interference and knowledge transfer in learning causal models in the context of continuous streams of data
- Effective ways for the causal structure to aid in leveraging the accumulated knowledge of a continual learner
- Leveraging causal tools to interpret distributional shifts in continual learning
- Next-generation benchmarks that go beyond re-purposing of existing datasets to adequately support the above items and further essential research questions towards a symbiosis of continual learning and causality

Overall, the call has resulted in 13 accepted papers, featuring an incredible breadth of topics. Accepted papers were able to concisely and clearly articulate general visions on linking fields, propose tangible mechanisms, highlight prospects for various applications, and already suggest first benchmarks. These papers have been collected in Volume 208 of the Proceedings of Machine Learning (PMLR), for which this retrospective serves as the foreword. All accepted papers were extended an invitation to participate in a poster session at the inaugural bridge event. Additionally, a limited number of papers with an outstanding degree of clarity were selected to present their perspectives in contributed talks.

2.2. Bridge schedule at a glance

Materials available at: https://www.continualcausality.org

Tutorials

Putting the *Continual* in Continual Causality
Keiland Cooper and Martin Mundt

Putting the *Causality* in Continual Causality
Devendra Singh Dhami and Adele Ribeiro

Software Labs

Avalanche: An End-to-End Library for Continual Learning
Antonio Carta

DoWhy: Causal Inference in Python
Peter Götz and Patrick Blöbaum

Vision Talks

The Neuro-Symbolic Continuum between Language, Knowledge and Reasoning
Yejin Choi

Continual Learning with Real-World Impact: Beyond Catastrophic Forgetting
Christopher Kanan

Going beyond the here and now: Counterfactual simulation in human cognition
Tobias Gerstenberg

Memory, Invariance and Reasoning: Pillars of the Causal-Continual Bridge
Vineeth Balasubramanian

Panel

Challenges and Opportunities of Continual Causality
Panelists: Christopher Kanan, Vineeth Balasubramanian, Dhireesha Kudithipudi, Moritz Grosse-Wentrup
Moderator: James Smith

Contributed Talks

Modeling Uplift from Observational Time-Series in Continual Scenarios
Sanghyun Kim, Jungwon Choi, NamHee Kim, Jaesung Ryu, Juho Lee

From IID to the Independent Mechanisms assumption in continual learning
Oleksiy Ostapenko, Alexandre Lacoste, Laurent Charlin

Never Ending Reasoning and Learning: Opportunities and Challenges
Sriraam Natarajan, Kristian Kersting

Towards Causal Replay for Knowledge Rehearsal in Continual Learning
Nikhil Churamani, Jiaee Cheong, Sinan Kalkan, Hatice Gunes

3. Discussion Insights and the Way Forward

The Continual Causality bridge featured several avenues for interactive discussion, motivated by the breadth of original visions and stemming from an expectation that a first effort is required to strive towards an initial consensus on the next steps. For this purpose, vision talks have summarized four fundamentally different interdisciplinary approaches to pinpoint relevant factors for continual causality. A panel, featuring a subset of vision speakers and further experts in the fields of continual learning and/or causality, has then continued to discuss select key questions and practical considerations. Ultimately, following contributed talks from exceptionally well-articulated paper submissions, the interactive discussion has been broadened to the full audience and virtual attendees through the formation of breakout groups. Here, each group has been balanced to be equally comprised of researchers attributing their research focus more towards continual learning or causality respectively. Each group has been confronted with the challenging task of narrowing down their impressions to the formulation of only *one* crucial challenge and direction, which they deem particularly worthy of immediate pursuit. We summarize the essence of these three discussion paths in the following.

3.1. Vision Talks

We briefly distill the essence of the vision talks and their impact on bridging continual causality, however, would like to defer the reader to the full materials available online to let appropriate credit remain with the speakers and the citations within their talks. Beyond a large amount of technical in-depth content, the four ensuing highlights illustrate the diversity aggregated by the bridge program.

In summary, an argument on the crucial role of knowledge and reasoning (to be integrated into continual learning), can perhaps best be attributed to directly paraphrasing an example: *"Birds in a vacuum, even when supplied with oxygen, cannot fly is something no one has likely ever said to you, but you nevertheless know it's true"* (Yejin Choi). This "dark matter" of language models, commonsense reasoning, remains a key open question in the development of language models and more generally in AI (Choi, 2022). A complementary perspective has nicely supplemented this fundamental insight by arguing that humans are capable, as experimentally evidenced, of counterfactual simulation. Correspondingly, our prospective causal continual learning should be able to mimic this behavior, by equipping them with mechanisms that can simulate what could happen. Stimulated by current continual learning practice, a position has been presented that benchmarks need to be reconsidered in continual learning, and more specifically, that computation and memory need to be key design elements. Intuitively, continual learning should naturally lend itself to various real-world applications due to these factors, yet academic papers seem to still largely focus on constructed examples. Lastly, a perspective on how to integrate causal mechanisms into the three pillars of alleviating catastrophic forgetting in continual learning: rehearsal, regularization, and dynamic architectures (Hadsell et al., 2020), has been illuminated. For instance, commonly used biologically inspired memory buffers (Hayes et al., 2021) could be extended to encompass examples that capture causal structure, continuous model regularization by distilling knowledge that respects causal properties (Hu et al., 2021; Kancheti et al., 2022), or linking dynamically growing architectures to causal discovery.

3.2. Panel Discussion

Practicality matters: is forgetting itself "solved"?

Continual learning algorithms have performed well using traditional approaches, such as memory buffers, to solve standard benchmarks. As the continual learning field moves away from its signpost in tackling catastrophic forgetting and into other field-related issues and domains, the integration of causal modeling may aid new and growing sub-fields, such as continual one-shot learning, learning composable features, or transfer learning. With these challenges, standard approaches to overcome interference may prove to be wanting, and may present a new opportunity for the integration of novel models to solve novel challenges. To get there, potential gaps in performance and knowledge will become apparent through the development of new benchmarks which holistically test models, for example, by prioritizing that continual learning algorithms should be resource aware by design.

Should causality imply being "actionable"?

Although causality in machine learning has grown exponentially in recent years with stupendous success in simple and synthetic environments, the question of it being useful in complex real world problems remains open. Concurrently, another critical open question is that of actionability. Whereas learning the underlying causal structure for a given problem is interesting, is it also profitable beyond a mere understanding if this learned causal structure cannot be acted upon? Continual causal models may naturally fit into this perspective, as they should focus on meaningful interventions and consecutive actions that have either led to a shift in data or will respectively induce changes. This can answer the question of actionability, since continual causality models should learn to act upon the learned causal structure in a dynamic fashion and could thus move the goalpost of the underlying models to be both more introspective and more retrospective.

Should ecological sustainability be at the forefront?

The successes of attaining top performance on a host of benchmarks by scaling model parameters has become a rapidly growing research area. Yet, these performance gains come with ecological and economical costs. To mitigate these shortcomings, as new models are developed in continual learning, causality, or their intersection, model and inference efficiency should not be an afterthought of development. Continual learning itself provides a solution to the issue of costs due to model retraining, by providing the tools needed to update existing models with new data without retraining *ab initio*. However, we should also acknowledge that our developments are often hardware specific or that particular made advances are at the very least more tailored towards the latter than some of the presently less popular hypothesized mechanisms. A symbiotic development with hardware, or respective exploration of alternate compute platforms beyond graphical processing units (GPUs), may be crucial to the evolution of continual learning, causality, and their interface. Additionally, improvements to benchmarks and datasets themselves may help models scale sustainably, through the development of smaller datasets which capture the main principals and advantages that larger datasets provide.

3.3. Interactive Breakout Discussions

Shareable knowledge and similarity in graphs

Traditionally, causal graphs are formed across a single dataset or domain. Yet, navigating real world tasks often would be aided by (or requires) leveraging learned causal structure from experience. To mitigate this issue, a continual learning agent may directly benefit from sharing its causal knowledge across tasks. This approach offers an immediate relationship between both continual learning and causal models, and exploits the similarity across many tasks. Learned knowledge may be shared forward in time, leveraging its past knowledge to solve similar or even novel tasks, and/or backwards in time, leveraging new data to validate, update, or refute previously learned causal associations. Across multiple tasks, a continual learner will in effect "meta-learn" a single or a core set of shared causal graph(s), refining core causal assumptions and facilitating performance across settings.

"Side-by-side" problems and benchmarks

An additional direction of further study is to look for existing areas where continual learning and causal discovery are already aligned. This may be through a shared structure of the problem at hand, or through shared domains. For example, medical data is already a topic of great research interest in both fields. In causal modeling and discovery, through a host of datasets to derive key variables and their relationships, and in continual learning, through learning models when access to the data is limited due to access or privacy and thus models are unable to be retrained. As new benchmarks to assess the progress of integrating the fields are developed, synergies should be identified, whereby leveraging areas where there is existing research, interest, and knowledge across both fields may prove to facilitate collaboration, interpretation, and integration.

Causal discovery and task ordering/curricula

Often, continual learning agents are agnostic to the order in which tasks are learned or there is little forethought into how specific knowledge from one task may directly affect (beneficial or otherwise) downstream learning. Causal discovery may aid this issue: by examining the structure of the tasks themselves, a continual learning researcher (or the agent itself) may have the tools to answer whether an optimal ordering exists for a sequence of tasks. From the standpoint of an agent, this approach may provide an automated method, based on explicit criteria drawn from existing continual learning metrics, to derive its own curricula, and to decide *which* information to learn *when*. From a research standpoint, this approach may help elucidate key principles of not only how multiple datasets are to be learned continually, but to perhaps provide a framework to answer a more stimulating question of *what* information needs to be extracted and retained from a prior task.

4. Conclusions

Whereas the differences between causality and continual learning may appear to be a gap too large to bridge, through the ideas, discussions, and collaborations of researchers across both fields, it is our conclusion that the possible benefits garnered by investigation into both field's intersection show mutually beneficial potential. For this, we are grateful to those who submitted manuscripts, the invited speakers, AAAI-23 organizers, and all attendees of the Continual Causality bridge program. We look forward to future discussions, collaborations, and research with this newfound community.

References

E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113:7345–7352, 2016.

Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3), 2018.

Yejin Choi. The curious case of commonsense intelligence. *Daedalus*, 151(2):139–155, 2022.

Zhixuan Chu, Stephen Rathbun, and Sheng Li. Continual lifelong causal effect inference with real world evidence. 2020.

Zachary Davis, Neil Bramley, Bob Rehder, and Todd M Gureckis. A causal model approach to dynamic control. In *CogSci*, 2018.

Tianwei Gong, Tobias Gerstenberg, Ralf Mayrhofer, and Neil Bramley. Active causal structure learning in continuous time. 2022.

Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12):1028–1040, 2020.

Tyler L. Hayes, Giri P. Krishnan, Maxim Bazhenov, Hava T. Siegelmann, Terrence J. Sejnowski, and Christopher Kanan. Replay in deep learning: Current approaches and missing biological elements. *Neural Computation*, 33(11):2908–2950, 2021.

Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3957–3966, 2021.

Khurram Javed, Martha White, and Yoshua Bengio. Learning causal models online. *arXiv preprint arXiv:2006.07461*, 2020.

Sai Srinivas Kancheti, Abbavaram Gowtham Reddy, Vineeth N Balasubramanian, and Amit Sharma. Matching learned causal effects of neural networks with domain priors. In *International Conference on Machine Learning*, pages 10676–10696. PMLR, 2022.

Dhireesha Kudithipudi, Mario Aguilar-Simon, Jonathan Babb, Maxim Bazhenov, et al. Biological underpinnings for lifelong learning machines. *Nature Machine Intelligence*, 4: 196–210, 2022.

Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.

S. Lee, J. Correa, and E. Bareinboim. Generalized transportability: Synthesis of experiments from heterogeneous domains. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, New York, NY, 2020. AAAI Press.

Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L. Hayes, et al. Avalanche: an end-to-end library for continual learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 3600–3610, 2021.

Sara Magliacane, Thijs Van Ommen, Tom Claassen, Stephan Bongers, Philip Versteeg, and Joris M Mooij. Domain adaptation by using causal inference to predict invariant conditional distributions. *Advances in neural information processing systems*, 31, 2018.

Martin Mundt, Steven Lang, Quentin Delfosse, and Kristian Kersting. Cleva-compass: A continual learning evaluation assessment compass to promote research transparency and comparability. *In International Conference on Learning Representations*, 2022.

J. Pearl and E. Bareinboim. External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595, Nov 2014.

Judea Pearl. *Causality.* Cambridge university press, 2009.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms.* The MIT Press, 2017.

Takeshi Teshima, Issei Sato, and Masashi Sugiyama. Few-shot domain adaptation by causal mechanism transfer. In *International Conference on Machine Learning*, pages 9458–9469. PMLR, 2020.

J Tian. Identifying dynamic sequential plans. In *Proceedings of the Twenty-fourth Conference on Uncertainty in Artificial Intelligence (UAI-08)*, page 554–561, Corvallis, Oregon, 2008. AUAI Press.