COST
EUROPEAN COOPERATION
IN SCIENCE & TECHNOLOGY

Short-Term Scientific Mission Grant
- APPLICATION FORM[1] -

Action number: CA20111

Applicant name: Mario Román

Details of the STSM

Title: A Type Theory for Exact and Continuous Bayesian Observations

Start and end date: 20/05/2024 to 31/05/2024

Goals of the STSM

Purpose and summary of the STSM.

This STSM aims to find a type theory for exact and real-valued Bayesian observations [9,10]. Multiple probabilistic programming languages have some form of "observe" statement [12]; observing something means updating our priors, according to Bayes' update, on the consideration of that event. Although there is much work on the denotational semantics of such a statement [13], formal reasoning on it remains elusive. This STSM aims to develop a synthetic type theory that allows reasoning with an "observe" statement.

Observations will be exact [11], meaning that it is not possible to declare to have observed a distribution but only a single element. This poses some problems when developing denotational semantics categorically, and it is easier to express by employing the difference between values and computations of call-by-push-value [6,7,8].

Observations will be continuous, meaning that we can observe a real number (or, more generally, a zero-measure result) as the output of a process. Traditionally, this complicates reasoning; however, our recent work on partial Markov categories [1] has proven that a syntactic procedure can reduce these continuous exact observations to computations with conditionals in a Markov category [2,3,4]. This novel ingredient is most interesting when included in a type theory: instead of considering a separate theory for the continuous case, we can develop a theory for both discrete and continuous conditionals.

The second goal of this STMS is to combine this development with our ongoing work on Markov bicategories [1,5] to allow reasoning about bounds and observations within a single unified type theory for probabilistic uncertainty.

References

[1] This form is part of the application for a grant to visit a host organisation located in a different country than the country of affiliation. It is submitted to the COST Action MC via-e-COST. The Grant Awarding Coordinator coordinates the evaluation on behalf of the Action MC and informs the Grant Holder of the result of the evaluation for issuing the Grant Letter.

[1] Di Lavore, Elena and Mario Román, "Evidential Decision Theory via Partial Markov Categories," *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, Boston, MA, USA, 2023, pp. 1-14, doi: 10.1109/LICS56636.2023.10175776.

[2] Fritz, Tobias. "A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics." *Advances in Mathematics* 370 (2020): 107239.

[3] Perrone, Paolo. "Markov categories and entropy." *IEEE Transactions on Information Theory* (2023).

[4] Fritz, Tobias, and Wendong Liang. "Free gs-monoidal categories and free Markov categories." *Applied Categorical Structures* 31.2 (2023): 21.

[5] Paquet, Hugo, and Philip Saville. "Effectful semantics in 2-dimensional categories: premonoidal and Freyd bicategories." *arXiv preprint arXiv:2312.14964* (2023).

[6] Paul Blain Levy, Call-by-push-value. Dissertation. Queen Mary University of London, UK, 2001.

[7] Chris Heunen, Bart Jacobs: Arrows, like Monads, are Monoids. MFPS 2006: 219-236.

[8] Paul Blain Levy, John Power, Hayo Thielecke: Modelling environments in call-by-value programming languages. Inf. Comput. 185(2): 182-210 (2003).

[9] Eliezer Yudkowsky, Nate Soares: Functional Decision Theory: A New Theory of Instrumental Rationality. CoRR abs/1710.05060 (2017).

[10] Arif Ahmed. Evidential Decision Theory. Cambridge University Press.

[11] Stein, Dario, and Sam Staton. "Probabilistic Programming with Exact Conditions." *Journal of the ACM* 71.1 (2024): 1-53.

[12] Staton, Sam. "Commutative semantics for probabilistic programming." *Programming Languages and Systems: 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22–29, 2017, Proceedings 26*. Springer Berlin Heidelberg, 2017.

[13] Heunen, C., Kammar, O., Staton, S., Moss, S., Vákár, M., Scibior, A., & Yang, H. (2018, January). The semantic structure of quasi-Borel spaces. In *PPS Workshop on Probabilistic Programming Semantics*.

Working Plan

Description of the work to be carried out by the applicant.

Practical aspects: The visitor (Mario Román) and the host (Elena Di Lavore) will work together, building on previous joint work [1]. A five-day allows time for focused, in-person collaboration, more lengthy in-depth discussions, and contrasting ideas with the rest of the Computer Science department in Pisa. This week will allow more intense work and coordination, making it easier to coordinate later towards the complete text and journal submission. It may be possible for the department to cover part of the accommodation costs.

Technical aspects: Exact observations are a common construct of probabilistic programming languages. However, a synthetic theory of exact observations has been missing until recently, when Stein and Staton [11] proposed a first categorical model of exact observations; our previous work has interpreted exact observations as partial structure in Markov categories [1]. In the same sense that Markov categories are a synthetic theory of probability, partial Markov categories give a synthetic theory of constrained probability that accounts for observations or Bayesian updates.

We plan to work explicitly towards the type theory (and the corresponding proof theory) that can be derived from this categorical structure. A synthetic probability theory could allow us to develop a theory of epistemology on a relatively short number of axioms that can be used to prove theorems about reasoning systems. We plan to investigate what is the internal language of partial Markov categories and to which kind of systems it can be applied.

In a previous STSM, we investigated the 2-categorical structure of partial Markov categories; we encountered some encouraging preliminary results suggesting it could be used to prove bounds on the space of possible observations and updates. However, one of the main limitations of this approach is that it seemed to work only for discrete cases – those that, classically, would not involve observations of measure zero. We plan to explicitly address this limitation, applying one of the main results of our previous work on partial Markov categories to the computation of exact observations [1,11].

Expected outputs and contribution to the Action MoU objectives and deliverables.

Main expected results and their contribution to the progress towards the Action objectives (either research coordination and/or capacity building objectives) and deliverables.

(max.500 words)

This research aligns with WG6 ("Type theories"). We expect a type theory for exact observations on partial Markov categories to fall within D4 ("Definition of a mathematical framework for modular reasoning about type theories and their extensions"); potentially, we could draft a plan for an implementation of this type theory, which would fall within D15 ("Prototype implementation of the mathematical framework, with basic user interface, user documentation and gallery of examples of type theories"). The new syntax would help bring Research Coordination Objectives 2 and 3 to the probabilistic programming field ("Promote the output of detailed, checkable proofs from automated theorem provers." and "Make techniques for program verification more effective and more accessible to all stakeholders"), which has traditionally not been concerned with formal verification and theorem proving.

We will finish our journal submission on "Partial Markov Categories" and expect to start preparing a submission highlighting the internal language of these categories. We aim to publish this work in a relevant theoretical computer science journal (e.g., LMCS, TheoretiCS).

Additionally, we would make progress towards the following Capacity Building Objectives: Objective 1 ("Bring together members of the different communities working on proofs in Europe") and Objective 3 ("Create an excellent and inclusive network of researchers in Europe with lasting collaboration beyond the lifetime of the action"), by contributing to a lasting collaboration between Pisa and Oxford; and Objective 4 ("Ease access to formal verification techniques in education and other areas of science"), by elaborating formal verification techniques for probabilistic programming.