▸ FRANCESCO A. GENCO, *Probabilistic computation and trust through the lens of typed λ-calculus* (a joint work with GIUSEPPE PRIMIERO).
LUCI Group, Department of Philosophy, University of Milan, Via Festa del Perdono 7, Milan, Italy.
E-mail: `francesco.genco@unimi.it`.

Computational systems evolve at a vertiginous speed and our notion of computation evolves with them. In recent years, the theory of computation embraced the study of systems that lie certainly outside its traditional borders. Probabilistic programming languages are a most prominent example of this phenomenon. The broadening of our notion of computation does not only force us to consider new mechanisms and systems, but also to rethink, in some cases, the role of some fundamental and well-established notions. The traditional notion of *correctness*, for instance, does not fare well at all with respect to probabilistic programs. A typical probabilistic program, indeed, cannot be said to compute *the* correct result. In spite of this, we often have quite strong expectations about the frequency with which it should return different outputs. Hence, instead of requiring programs to be *correct*, we can appeal to a weaker notion and require them to be *trustworthy*, where by *trustworthy* we mean that they yield outputs as determined by probability distributions that model their expected behaviour.

We present a computational framework that precisely formalises this intuition. In order to do so, we define a typed λ-calculus that features operators for conducting experiments on probabilistic programs and for evaluating whether the frequency of their outputs is compliant with target probability distributions. We then discuss the fundamental computational properties of the calculus: subject reduction, progress and termination.

The analysis of the notion of trust provided by the calculus is essentially pragmatic in nature and, coherently, heavily depends on the contingencies of the particular program execution under consideration. In order to also provide a more abstract, atemporal characterisation of the investigated phenomena, we introduce a notion of *confidence* which does not depend on the development of any particular execution of a program but only on its static definition. This enables us to prove—by an application of the strong law of large numbers—that the calculus and the notion of trust formalised in it globally behave as expected with respect to the basic tenets of probability theory.

Keywords. Trust, probabilistic computation, type theory, lambda-calculus.