"Semantic Pollution of Proof Systems"

Robin Martinot

Abstract

Proof systems and model-theoretic semantics provide different ways of proving results about logics, and soundness and completeness proofs reveal an intrinsic connection between these methods. However, Avron (1996) writes that a requirement of a 'good' proof system is that it should be "independent from any particular semantics". This has become known as syntactic purity of a proof system, as opposed to a semantically polluted one. The value of soundness and completeness proofs seems to come from a certain independence that the syntactic side has from the semantic side. If a proof system is semantically polluted, this may take away from its "proof-theoretical nature and the expected generality" (Avron, 1996). Labeled proof calculi are a standard example of semantically polluted systems, for directly internalizing Kripke semantics in the proof system (see e.g. Poggiolesi and Restall, 2012), but other examples are spread across the literature, including for instance semantic sequents and tableaux and internalized forcing sequents (Poggiolesi, 2010), or the inclusion of neighborhood semantics into the proof system (Negri, 2017).

The goal of this talk is to provide a better conceptual characterization of what semantic pollution is, and to provide and compare formal ways of telling when a proof system is semantically polluted or not. This contributes to a better characterization of what a 'good' proof system can be, and encourages a more nuanced understanding of the distinction between syntax and semantics. The literature distinguishes between a strong and a weak definition of syntactic purity. Strong syntactic purity occurs when a proof system is "independent of any particular semantics" (Avron, 1996). This includes the idea that "one should not be able to guess, just from the form of the structures which are used, the intended semantic of a given proof system". Weak syntactic purity, on the other hand, says that a sequent calculus cannot make use of 'explicit semantic elements' (Poggiolesi, 2010). Poggiolesi argues that strong syntactic purity is too strong, since it implies that basic propositional sequent calculi already have to be declared semantically polluted. Thus, she adopts weak syntactic purity, where she defines a 'semantic element' as an untranslatable ingredient of a sequent (as compared to a usual original language). This rules out, for example, expressions like xRy in labeled calculi, that explicitly incorporate the notion of possible worlds and the Kripke accessibility relation.

In this talk, we first discuss several informal conceptions of semantic pollution. We suggest that strong syntactic purity might not be too strong after all, by investigating the idea that you can 'guess' the intended semantics just from the rules of a proof system. As Hacking (1979) notes, independently deducing the semantic meaning of inference rules relies on strong semantic assumptions, that limit any real guessing process. Thus, we explore weaker but also more formal understandings of 'guessing' a semantics. We also provide a conception of a 'semantic element' so that it, in a sufficient way, has "an evident connection with truth or warranted assertibility" (Dummett, 1998).

Based on these ideas, we move to possible ways to formally characterize semantic pollution. For weak syntactic purity, we discourage the idea that translatability is decisive in the formal description of a semantic element. Instead, we aim to spell out requirements on the formal language in a proof system in order to exclude semantic elements, which also helps us understand better why these elements are excluded. Among such requirements is the idea that symbols should be able to intermingle

sufficiently, and that they should not be able to only say something about the models. We also discuss the view that the level of 'explicitness' of representation of semantic elements is important. Namely, Poggiolesi and Restall (2012) note that elements from Kripke semantics are treated explicitly in labeled systems, but are made implicit in tree-hypersequent systems (reducing the level of semantic pollution). Read (2015) objects that even in tree-hypersequent systems, "the content is still there". We argue that the particular presentation of content does indeed matter for semantic pollution. For strong syntactic purity, finally, we focus on possible formalizations of 'recognizing' semantics from a proof system. For example, the likeness of a proof line to its semantic definition of validity, the degree to which that a syntactic proof simulates a semantic proof, or the way that proof rules determine the semantics of the logical connectives they define, might relate to semantic pollution.

We conclude by considering the implications of this work for various proof systems, and we reflect on whether our measures of semantic pollution seem to interact with any other philosophical properties.