

Identifiers should be simple: just something that you can use to identify a resource. But they turn out to be surprisingly complex with a very large number of ways of doing it wrong. We can split most identifier schemes into two broad camps: with semantics and semantics-free. Over the last decade, semantics-free numerical identifiers have largely been considered the best, and often seen as best practice for use in ontologies. But they also come with a lot of problems. I will describe some of these issues, and put forward a proposal that I am working on for better identifiers.