

The vast majority of current multi-model querying solutions require the user to have intimate knowledge of the specific models involved. There exists a single approach for truly unified multi-model querying, but this approach is not practically usable for most users due to its complexity. In this thesis we present MMQL, a multi-model query language based on category theory, which was designed using SPARQL as a basis. Using MMQL, users can formulate multi-model, multi-database queries without needing to know about the way the data is stored. We also present our proposal for the implementation of MMQL, including the required supporting algorithms. To verify the validity of our proposal, we built the proof-of-concept tool MM-quecat, an implementation of basic MMQL concepts. We then evaluated MM-quecat in a scenario involving PostgreSQL and MongoDB, querying both databases with a single MMQL query. As we present one of the first ever approaches for unified multi-model querying, we also analyze the weaknesses and limitations of the proposed approach, opening the door for future iterations and improvements.