

Table-to-text generation, a subtask of Natural Language Generation (NLG), involves generating coherent and faithful text from structured tables. Modern approaches to this task have overcome language fluency problem but still struggle with generation fidelity, especially in cases where reasoning is required. This thesis focuses on Logical Table-to-Text (LT2T) generation, addressing fidelity, as all previous studies do, but also examining a yet underexplored problem of processing potentially long tables. We investigate the possibility to select the relevant content without using the entire table, as well as the role of symbolic reasoning in enhancing statement fidelity. Our pipeline-based system generates a template of an intermediate logical form from table aggregation and fills the template using a symbolic approach. It outperforms previous works in fidelity and produces natural outputs, as assessed through automatic metrics and human evaluation.