

METHOD AND APPARATUS TO FACILITATE BUILDING A LARGE LANGUAGE MODEL PIPELINE

Abstract of the Disclosure

A textual description of an issue pertaining to at least part of an apparatus is received as input. A data store is accessed and a plurality of knowledge documents that corresponds to the input is retrieved. A language generation prompt is then generated, as a function of both the input and the plurality of knowledge documents and output to a task-specific decoder that generates a candidate recommendation to address the aforementioned issue. That candidate recommendation is output to at least one human reviewer who reviews the candidate recommendation as a function of the plurality of knowledge documents and who then provides a corresponding human-validated recommendation to address the issue. The task-specific decoder can be retrained using the human-validation recommendation coupled with the corresponding textual description input.

[Figure 1]