Supplementary Note 2: Question answering templates.

This section presents the question templates used to construct the question-answering dataset. Each template consists of three components: QUESTION, SLOTS, and EVIDENCE PATH. SLOTS defines the ontology-constrained entity types to be instantiated.

For brevity, Templates T01–T12 are expressed at the superclass level “Culture-loaded Word”. In actual instantiation, this slot may correspond to any of its three subtypes—character, word, or sentence—depending on the granularity of the queried expression. To make this explicit, we provide subtype-specific instantiation examples after the core template set.

EVIDENCE PATH specifies the expected ontology-level path shape and the typical retrieval pattern along ontology relationships. In the framework, it first acts as a structural prior in Stage 2 to constrain admissible entity-type paths; paths inconsistent with the extracted anchor elements, target element, or maximum search depth are pruned before entity-path instantiation. The retained instantiated paths are then semantically scored by the node scoring model in Stage 3, and only the top-k highest-scoring paths are preserved for answer generation. Thus, EVIDENCE PATH constrains the structurally admissible retrieval space, while the node scoring model determines which of the retained candidate paths are selected as final evidence. Operationally, a candidate path enters Stage 3 only if its ontology-level path is compatible with the extracted anchor elements, target element, and maximum search depth; the node scoring model never expands the search space, but only reorders the structurally admissible candidate paths retained after Stage 2.

The template set spans cases from 1-hop to 4-hop, supporting the unified construction, analysis, and evaluation of questions with varying reasoning depths.