SN-YE: An Ontology-Powered Tool for Exploring Surya Namaskar Yoga

Bharath Muppasani[1], **Mansi Dodiya**[2], **Kishan Kumar**[2],
Raghava Mutharaju[3], **Biplav Srivastava**[1], **Hari Prabhat Gupta**[2]
[1]University of South Carolina, USA
[2]Indian Institute of Technology (BHU), Varanasi, India
[3]Mehta Family School of Data Science and AI, IIT Palakkad, India
bharath@email.sc.edu, mansibhavikbhaidodiya.cse23@itbhu.ac.in, kishan.jrf.cse25@itbhu.ac.in, raghava@iitpkd.ac.in,
biplav.s@sc.edu, hariprabhat.cse@iitbhu.ac.in

Abstract

Yoga is practised worldwide for health, but the knowledge behind it is scattered across instructional books, tradition-specific manuals, and videos that differ in vocabulary, detail, and language. So a learner has no easy way to ask a precise question. Surya Namaskar (SN), a widely practised yoga sequence, comprises 12 poses drawn from 9 unique asanas (postures) and is practised in four variants represented in our knowledge base. We present **SN-YE**, the Surya Namaskar Yoga Explorer: an ontology-powered tool that answers such questions from a knowledge graph rather than from free text. It runs contextual queries in the SPARQL Protocol and Resource Description Framework Query Language (SPARQL) over the Surya Namaskar Yoga Ontology (SN-YO) and returns retrieved facts alongside illustrations from the Common Yoga Protocol (CYP). A *user view* offers pre-canned competency questions answered in any of 28 languages; a *developer view* lets a large language model (LLM) translate free-form questions into SPARQL, or lets SPARQL be written directly. The model only plans the query, and every fact comes from local execution over the ontology. SN-YE is a static web page that runs entirely client-side, extends to new asanas and Surya Namaskar variants by adding ontology modules, and can be embedded into any web application.

Demo Video: https://tinyurl.com/ai4s-sn-ye
Explorer Interface: https://ai4society.github.io/sn-ye/

1 Introduction

Surya Namaskar (Sun Salutation) is one of the most widely practised Yoga routines: an ordered sequence of postures, each coordinated with a breathing phase and, in some traditions, with a mantra and a chakra. It is studied for its effects on strength and endurance (Bhutkar et al. 2011). Yet the knowledge a practitioner needs is hard to reach. It lives in instructional documents, lineage-specific books, and videos; traditions order and name postures differently; the same posture appears under Sanskrit, transliterated, and translated names; and corrective guidance is buried in prose. Artificial intelligence work on Yoga has largely targeted perception, such as fine-grained pose classification (Verma et al. 2020), rather than making practice knowledge queryable.

Our key idea is to organise this knowledge in an ontology and place *contextual query tools* on top of it, with query translation handled by a large language model (LLM). The separation matters: LLMs map informal questions onto formal query intents well, but are poor sources of safety-relevant facts. In SN-YE, the Surya Namaskar Yoga Explorer, the model never answers; it decides *what to ask*, and the ontology decides *what is true*, following ontology-driven tools built for planning knowledge (Muppasani et al. 2024, 2026).

Copyright © 2027, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

2 SN-YO: The Knowledge Source

SN-YE is powered by the Surya Namaskar Yoga Ontology (SN-YO), a modular ontology in the Web Ontology Language (OWL) (Dodiya et al. 2026). Each module is published under a Persistent Uniform Resource Locator (PURL). Its central modelling decision separates an `Asana`, the identity of a posture, from a `Pose`, an *occurrence* of that asana at a position in a sequence. This lets Pranamasana appear at positions 1 and 12 of a round and be recognised as the same posture across traditions. Around this hub (Figure 1, inset), a pose links to its variant, asana, breathing phase, safety notes, and common errors, each error carrying a correction; `hasNextPose`, `repeatsPose`, and `hasInversePose` capture order, symmetric recurrence, and left/right pairing. The bundled version covers four variants (Sivananda, Krishnamacharya, Bihar School of Yoga, and Vivekananda Kendra) over 49 pose occurrences and 33 asanas, with 35 pose–error links, 34 corrections, and 24 asanas linked to pages of the Common Yoga Protocol (CYP) (Ministry of Ayush, Government of India 2022). Labels are given in 28 languages: the six official United Nations (UN) languages and the 22 languages of the Eighth Schedule of the Constitution of India.

3 The SN-YE Explorer

SN-YE is a static web application that parses the SN-YO OWL file in the browser and evaluates SPARQL (Harris and Seaborne 2013) SELECT queries, including triple patterns, `OPTIONAL`, `VALUES`, `FILTER`/`LANGMATCHES`, `ORDER BY`, and `LIMIT`, with a local runner. There is no server, database, or endpoint to keep alive. Figure 1 shows the pipeline, which two views expose.

User view. This view presents ten pre-canned competency questions from the ontology's own evaluation set: sequence order, breathing and safety, mantra and chakra coverage, errors and corrections, repeats, laterality inverses, shared

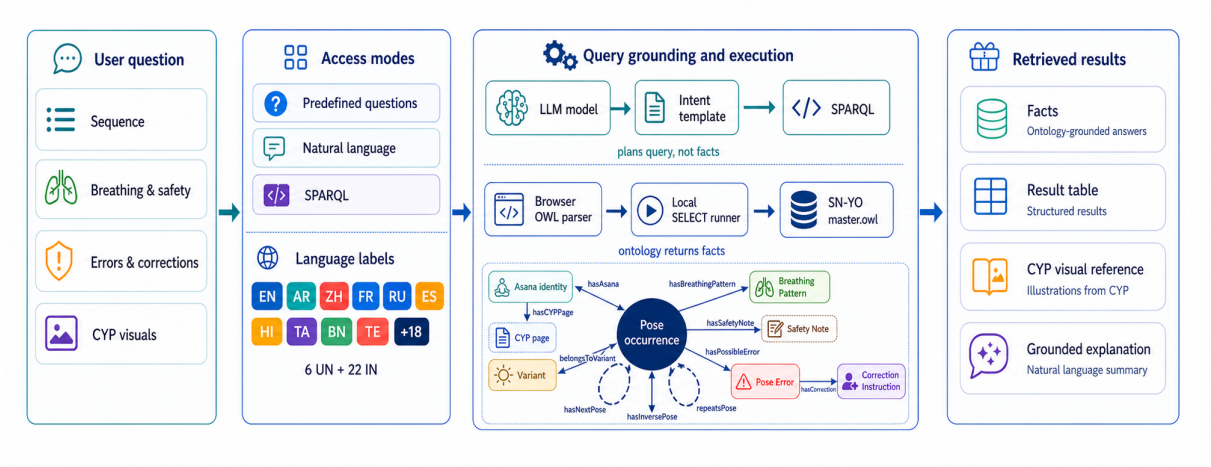


Figure 1: The SN-YE pipeline. A question entered in any of the three access modes becomes a SPARQL query; a language model, when involved, only plans that query by choosing an intent template, while every fact comes from local execution of the SELECT against the SN-YO graph in the browser. The inset shows the reduced SN-YO schema, in which a *pose occurrence* is the hub linking an asana identity, a variant, a breathing pattern, safety notes, and error and correction knowledge, while pose-to-pose relations capture ordering, repetition, and laterality inversion.

asanas across variants, asana equivalences, the variant catalogue, and visual references. Selecting one, runs its SPARQL template and renders a narrative, headline facts, and a result table whose columns follow the ontology: asana, breathing phase, laterality, support type, chakra, mantra, and safety note. A language selector re-runs the query with a different label language, so the same answer can be read in Hindi, Telugu, Arabic, or 25 others. Answers touching an asana with a linked CYP page also show it, anchoring an abstract result to an authoritative illustration.

Developer view. The developer view offers three workspace modes over the same graph. *Predefined* mode serves the ten curated questions, and each answer can be toggled to reveal the SPARQL template that produced it and the ontology data object it returned. *Natural-language* mode lets the user pick a model, supply a key, and type a free-form prompt: SN-YE sends the question with a serialised context of the loaded ontology, and the model must return a structured plan naming one of thirteen supported intents plus slots such as an asana or pose number. That plan is validated against a JSON schema and the loaded graph, compiled into SPARQL, and shown for review before it runs; an evidence-constrained second call then explains the retrieved rows. Because the model may not introduce schema terms, a question outside the intent set is reported as unsupported rather than answered. *SPARQL* mode provides an editor with the five SN-YO namespace prefixes, term hints, and worked examples, and a result panel reporting rows, bound variables, and the applied limit. The language selector applies in all three modes, which share the user view's execution path.

4 Demonstration

The demonstration runs on a laptop with the ontology loaded from disk, so it needs no network. In the user view, one can request the ordered base sequence and get twelve poses with their asanas, breathing phases, laterality, chakra, mantra, and safety notes, plus the five CYP pages that ground them; switching to Hindi re-renders the same rows in Devanagari with no new query, and asking which errors are documented per pose returns the 35 pose–error links with their corrections. In the developer view, one can ask the same questions in free-form English, inspect the planned SPARQL, run it, and hand-edit it to go beyond the pre-canned set. In free-form, one can ask an unsupported question (e.g., non-Yoga) and the system should be able to handle it as well (i.e., show no supported query).

5 Discussion and Conclusion

Two properties make SN-YE useful beyond the routine. It is **extensible**: coverage grows by adding OWL modules, not code. A new variant is a module of pose occurrences reusing existing asana identities, which cross-variant queries pick up automatically. It is also **embeddable**: as plain HyperText Markup Language (HTML), Cascading Style Sheets (CSS), and JavaScript with the ontology as a static asset, it drops into any website or wellness app and runs offline; only the natural-language mode needs an external service. Its main limitation is coverage: the intent set is fixed, which keeps the tool faithful but sends novel questions to the SPARQL editor; the multilingual labels are script-adapted renderings of Sanskrit-derived posture names rather than translations.

SN-YE shows that a practice with well-defined structure but a fragmented literature becomes explorable by combining a modest ontology with an LLM used only as a query planner: users get multilingual, visually grounded answers traceable to a query they can read, and developers get an auditable path from question to SPARQL.

References

Bhutkar, M. V.; Bhutkar, P. M.; Taware, G. B.; and Surdi, A. D. 2011. How Effective Is Sun Salutation in Improving Muscle Strength, General Body Endurance and Body Composition? *Asian Journal of Sports Medicine*, 2(4): 259–266.

Dodiya, M.; Kumar, K.; Muppasani, B.; Srivastava, B.; Mutharaju, R.; and Gupta, H. P. 2026. SN-YO: An Ontology for Modeling Surya Namaskar as a Sequence-Based Yoga Practice. In *Under review*.

Harris, S.; and Seaborne, A. 2013. SPARQL 1.1 Query Language. W3C Recommendation.

Ministry of Ayush, Government of India. 2022. Common Yoga Protocol. https://yoga.ayush.gov.in/.

Muppasani, B.; Dey, R.; Srivastava, B.; and Narayanan, V. 2026. OMEGA: An Ontology-Driven Tool for Explaining Multi-Agent Path Finding. In *Proceedings of the AAAI Conference on Artificial Intelligence (Demonstration Track)*.

Muppasani, B.; Gupta, N.; Pallagani, V.; Srivastava, B.; Mutharaju, R.; Huhns, M. N.; and Narayanan, V. 2024. Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications. In *Eighth International Conference on Data Science and Management of Data (CODS-COMAD'24)*.

Verma, M.; Kumawat, S.; Nakashima, Y.; and Raman, S. 2020. Yoga-82: A New Dataset for Fine-Grained Classification of Human Poses. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*.