



Tikos Reasoning Platform

The system creates two data assets for each ML/AI model: Cases and Context, key components for the Tikos case-based reasoning platform.

‘Cases’ is a proprietary data structure (1) - including activation path information for deep-neural networks (2) - which are optimised through information minimisation and serialised for efficient case indexing, searching, retrieval, matching, and adaptation (3). This process delivers log level monitoring and observability (4) for individual decision outputs for any model.

‘Context’ extends system capabilities from observation to explainability. Model features are combined with relevant domain information and represented in a knowledge graph, or other datastore (5). Matched or adapted Cases relating to individual model output decisions are then explained using the Context (6).

Innovations

Tikos is built on a family of proprietary formal methodologies, mathematical techniques and algorithms designed to work in concert to depoly the system at scale:

Sequential Collection: Data structure for knowledge serialisation & propagation

Synapses Logger: Process for capturing neural activation path information

ReduceBySQ: Distributed processing for Case search, retrieval and adaptation

Open Architecture

Tikos is intentionally agnostic to:

- Model class
- Developer framework
- Tooling
- Deployment infrastructure

And accessible through SaaS by APIs and SDK.

Expert in the Loop

Tikos can leverage insights of subject matter experts at key points during system build and operation:

- Design of Case data
- Domain information
- Context ontology
- Validation of outputs

