Hypertension and Total-Order Forward Decomposition Optimizations (Abstract Reprint)

Maurcio Ceclio Magnaguagno[1], **Felipe Meneguzzi**[1,2], **Lavindra de Silva**[3]

[1]Pontifical Catholic University of Rio Grande do Sul, Porto Alegre, Brazil
[2]University of Aberdeen, Aberdeen, UK
[3]Department of Engineering, University of Cambridge, Cambridge, UK

Abstract Reprint. This is an abstract reprint of the journal article by Magnaguagno, Meneguzzi, and de Silva (2025).

Abstract

Hierarchical Task Network (HTN) planners generate plans using a decomposition process with extra domain knowledge to guide search towards a planning task. Domain experts develop such domain knowledge through recipes of how to decompose higher level tasks, specifying which tasks can be decomposed and under what conditions. In most realistic domains, such recipes contain recursions, i.e., tasks that can be decomposed into other tasks that contain the original task. Such domains require that either the domain expert tailor such domain knowledge to the specific HTN planning algorithm, or an algorithm that can search efficiently using such domain knowledge. By leveraging a three-stage compiler design we can easily support more language descriptions and preprocessing optimizations that when chained can greatly improve runtime efficiency in such domains. In this paper we evaluate such optimizations with the HyperTensioN HTN planner, winner of the HTN IPC 2020 total-order track.

References

Magnaguagno, M. C.; Meneguzzi, F.; and de Silva, L. 2025. Hypertension and total-order forward decomposition optimizations. *Autonomous Agents and Multi-Agent Systems*, 39: 24.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.