Hapori Delfi

Patrick Ferber[1], Michael Katz[2], Jendrik Seipp[3], Silvan Sievers[1], Daniel Borrajo[4], Isabel Cenamor, Tomas de la Rosa, Fernando Fernandez-Rebollo[4], Carlos Linares López[4], Sergio Nuñez, Alberto Pozanco, Horst Samulowitz[2], Shirin Sohrabi[2]

[1] University of Basel, Switzerland
[2] IBM T.J. Watson Research Center, Yorktown Heights, USA
[3] Linköping University, Sweden
[4] Universidad Carlos III de Madrid, Spain

patrick.ferber@unibas.ch, michael.katz1@ibm.com, jendrik.seipp@liu.se, silvan.sievers@unibas.ch, dborrajo@ia.uc3m.es, icenamorg@gmail.com, tomdelarosa@gmail.com, ffernand@inf.uc3m.es, clinares@inf.uc3m.es, sergio.nunez@repsol.com, alberto.pozanco@gmail.com, samulowitz@us.ibm.com, shirin.sohrabi@gmail.com

Abstract

The cost-optimal planner *Delfi* has successfully participated in the International Planning Competition (IPC) 2018. Its success can be attributed to two main factors: the use of state-of-the-art cost-optimal planners in its portfolio and the ability to predict which of these planners is a good fit for a given planning task. Following that prior success, here we extend the set of planners in the portfolio. The learning methodology is adapted according to the prior work, applied now not only to cost-optimal, but also to agile and satisficing planning.

Introduction

The cost-optimal planner *Delfi* (Katz et al. 2018b) was rated first in the cost-optimal track of the International Planning Competition (IPC) 2018. It uses a so-called online portfolio approach (Cenamor, de la Rosa, and Fernández 2016; Sievers et al. 2019a) to overcome the limitations of any individual planner, predicting which out of the collection of planners will work well on the planning task at hand. That collection included 17 planners of mostly similar configurations, varying mostly in the heuristic used. The prediction was done with the help of a deep learning tool, specifically convolutional neural network (CNN) (LeCun, Bengio, and Hinton 2015), predicting whether a planner will solve a planning task, represented by an image, within the predefined time limit of 30 minutes. The image representation was obtained based on a structural representation of a planning task called *abstract structure graph* (ASG) (Sievers et al. 2019b), casting the graph as an adjacency matrix, condensing and turning into a grayscale image.

In this work, we construct a new, community based version of Delfi, which we now call *Hapori Delfi* [1]. We extend the collection of planners in the portfolio and adopt the best performing learning methodology and architecture according to the post-IPC 2018 investigation (Sievers et al. 2019a). Specifically, we discretize the time interval into three equal size intervals and predict whether the planner will solve the task within that time interval. We use the same image-based planning task representation as in the original *Delfi* and retrain the CNN for the new collection of planners. Additionally, we go beyond just cost-optimal planning, preparing versions of *Hapori Delfi* also for the agile and satisficing tracks. In the rest of the paper we describe the differences from the original *Delfi* for each of the tracks we participate in, specifically in the components used and the training data.

[1] Hapori is the Maori word for community.

Components and Training Data

As the pool of planners for our portfolios to choose from, we use all planners from the IPC 2018 and a selection of planners from IPC 2014. If an IPC 2018 planner is itself a portfolio, we use its component planners instead. We only consider each planner once. (Some IPC 2018 portfolios include planners that were also submitted separately and several portfolios included the same planners.)

Optimal Planners. For the optimal track, we exclude the planners MAPlan-1, MAPlan-2 and Meta-Search Planner because they use CPLEX, and Complementary1 because it may generate suboptimal solutions. Furthermore, the FDMS planners and Metis1 are covered by the Delfi portfolio already. This results in the following list of planners (or their components):

- Complementary2 (Franco, Lelis, and Barley 2018)
- components of DecStar (Gnad, Shleyfman, and Hoffmann 2018)
- components of Delfi (Delfi1 and Delfi2 have the same components; Katz et al., 2018b)
- Metis2 (Sievers and Katz 2018)
- Planning-PDBs (Moraru et al. 2018)
- Scorpion (Seipp 2018b)
- SymBA*1 (IPC 2014; Torralba et al., 2014)
- Symple-1 and Symple-2 (Speck, Geißer, and Mattmüller 2018)

Satisficing Planners. All planners participating in the IPC 2018 satisficing track also participated in the agile track (except for Fast Downward Stone Soup 2018), with an identical code base but possibly with different configurations. We thus only have one set of planners but multiple configurations for these two tracks. We exclude the Alien planner because we could not get it to run, and Freelunch-Doubly-Relaxed, FS-blind and FS-sim because they have a large number of dependencies which results in planner images too large to be included in our planner pool. Furthermore, IBaCoP-2018 and IBaCoP2-2018 use a large number of planners or portfolios of which newer and stronger versions participated in IPC 2018 as standalone planners, or which we failed to get to run, so we only cover the component planners Jasper, Madagascar, Mercury, and Probe. This results in the following list of planners (or their components):

- Cerberus and Cerberus-gl (Katz 2018)
- components of DecStar (Gnad, Shleyfman, and Hoffmann 2018)
- components of Fast Downward Remix (Seipp 2018a)
- components of Fast Downward Stone Soup 2018 (Seipp and Röger 2018)
- Jasper (IPC 2014; Xie, Müller, and Holte, 2014)
- Dual-BFWS, BFWS-preference, BFWS-polynomial and DFS$^+$ (Francès et al. 2018)
- Madagascar (IPC 2014; Rintanen, 2014)
- Mercury2014 (Katz and Hoffmann 2014)
- MERWIN (Katz et al. 2018a)
- OLCFF (Fickert and Hoffmann 2018)
- Probe (IPC 2014; Lipovetzky et al., 2014)
- Grey Planning configuration of Saarplan (Fickert et al., 2018; rest covered by DecStar)
- Symple-1 and Symple-2 (Speck, Geißer, and Mattmüller 2018)

Benchmarks and Runtimes. For training the portfolios, we use all tasks and domains from previous IPCs, from the Delfi training set (Katz et al. 2018b), and from the Autoscale 21.11 collection Torralba, Seipp, and Sievers (2021), leading to a set of 92 domains with 7330 tasks. We use Downward Lab (Seipp et al. 2017) to run all planners across all benchmarks on AMD EPYC 7742 2.25GHz processors, imposing a memory limit of 8 GiB and a time limit of 30 minutes for optimal planners and 5 minutes for satisficing and agile planners. For each run, we store its outcome (plan found, out of memory, out of time, task not supported by planner, unexpected error), the execution time, the maximum resident memory, and if the run found a plan, the plan length and plan cost. This data set is available online.[2] As training data for our optimal (respectively satisficing/agile) portfolios, we select from each domain the 30 tasks which are solved by the fewest optimal (or satisficing/agile) planners, which results in 1926 (optimal) and 2377 (satisficing/agile) tasks.

[2]https://github.com/ipc2023-classical/planner19/tree/latest/experiments/data/01-opt-planners-eval and https://github.com/ipc2023-classical/planner19/tree/latest/experiments/data/02-sat-planners-eval

Post-IPC Analysis

The IPC 2023 used 7 domains with 20 tasks each, resulting in a benchmark set of 140 planning tasks, for all three tracks. Each planner was limited to 30 minutes of CPU time and 8 GiB of memory.

In the optimal track, there were 22 competing planners. The objective was to optimally solve the tasks. The best planner solved 77 tasks, the blind baseline 50, the LM-cut baseline 34, and our portfolio only 41 tasks, ranking 12th. Unfortunately, our planner had some technical problems, such as selecting a component planner which did not support some PDDL features. Furthermore, it always selected the same component planner, namely Planning-PDBs, for all tasks. Most likely, the fact that we did not adapt the training setting (such as hyper parameters) compared to the previous version of Delfi to accomodate the new scenario lead to poor performance on the test set, i.e., the IPC 2023 benchmarks.

In the satisficing track, there were 22 competing planners. The objective was to find plans of high quality. The best planner achieved a summed score of 71.86, only closely beating the baseline LAMA with a score of 68.76, and our planner scored 32.54, ranking 18th. Here, too, our planner had some bugs that lead to selecting a planner by the wrong name or selecting a component planner which did not support some PDDL features. In the successful cases, our planner selected eleven out of all available component planners.

In the agile track, there were 22 competing planners. The objective was to find plans as quickly as possible. The best planner achieved a score of 40.25, closely below the baseline LAMA-first with a score 40.28, and our planner scored 10.83, ranking 18th. Even though the tasks used for the agile track are identical to those of the satisficing track, and the set of component planners available to our portfolios is identical for both tracks, our planner made a slightly different selection of planners due to not being deterministic, choosing twelve out of all available component planners, but facing the same technical problems as in the satisficing track.

We are actively working on fixing the bugs of our planners, and aim to do a thorough comparison of the Hapori portfolios in a subsequent journal article.

Acknowledgments

The success of a portfolio planner must be primarily attributed to the developers of the portfolio components. Therefore, we would like to express our gratitude to the numerous authors of the components on which our portfolios are based. We also thank Daniel Fišer and Florian Pommerening for organizing the competition and taking on the time-consuming task of running our numerous planner submissions.

References

Cenamor, I.; de la Rosa, T.; and Fernández, F. 2016. The IBaCoP Planning System: Instance-Based Configured Portfolios. *Journal of Artificial Intelligence Research*, 56: 657–691.

Fickert, M.; Gnad, D.; Speicher, P.; and Hoffmann, J. 2018. SaarPlan: Combining Saarland's Greatest Planning Techniques. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 11–16.

Fickert, M.; and Hoffmann, J. 2018. OLCFF: Online-Learning h^{CFF}. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 17–19.

Francès, G.; Geffner, H.; Lipovetzky, N.; and Ramiréz, M. 2018. Best-First Width Search in the IPC 2018: Complete, Simulated, and Polynomial Variants. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 23–27.

Franco, S.; Lelis, L. H. S.; and Barley, M. 2018. The Complementary2 Planner in the IPC 2018. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 32–36.

Gnad, D.; Shleyfman, A.; and Hoffmann, J. 2018. DecStar – STAR-topology DECoupled Search at its best. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 42–46.

Katz, M. 2018. Cerberus: Red-Black Heuristic for Planning Tasks with Conditional Effects Meets Novelty Heuristic and Enchanced Mutex Detection. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 47–51.

Katz, M.; and Hoffmann, J. 2014. Mercury Planner: Pushing the Limits of Partial Delete Relaxation. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 43–47.

Katz, M.; Lipovetzky, N.; Moshkovich, D.; and Tuisov, A. 2018a. MERWIN Planner: Mercury Enchanced With Novelty Heuristic. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 53–56.

Katz, M.; Sohrabi, S.; Samulowitz, H.; and Sievers, S. 2018b. Delfi: Online Planner Selection for Cost-Optimal Planning. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 57–64.

LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436–444.

Lipovetzky, N.; Ramirez, M.; Muise, C.; and Geffner, H. 2014. Width and Inference Based Planners: SIW, BFS(f), and PROBE. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 6–7.

Moraru, I.; Edelkamp, S.; Martinez, M.; and Franco, S. 2018. Planning-PDBs Planner. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 69–73.

Rintanen, J. 2014. Madagascar: Scalable Planning with SAT. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 66–70.

Seipp, J. 2018a. Fast Downward Remix. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 74–76.

Seipp, J. 2018b. Fast Downward Scorpion. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 77–79.

Seipp, J.; Pommerening, F.; Sievers, S.; and Helmert, M. 2017. Downward Lab. https://doi.org/10.5281/zenodo.790461.

Seipp, J.; and Röger, G. 2018. Fast Downward Stone Soup 2018. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 80–82.

Sievers, S.; and Katz, M. 2018. Metis 2018. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 83–84.

Sievers, S.; Katz, M.; Sohrabi, S.; Samulowitz, H.; and Ferber, P. 2019a. Deep Learning for Cost-Optimal Planning: Task-Dependent Planner Selection. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, 7715–7723. AAAI Press.

Sievers, S.; Röger, G.; Wehrle, M.; and Katz, M. 2019b. Theoretical Foundations for Structural Symmetries of Lifted PDDL Tasks. In Lipovetzky, N.; Onaindia, E.; and Smith, D. E., eds., *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*, 446–454. AAAI Press.

Speck, D.; Geißer, F.; and Mattmüller, R. 2018. SYMPLE: Symbolic Planning based on EVMDDs. In *Ninth International Planning Competition (IPC-9): Planner Abstracts*, 91–94.

Torralba, Á.; Alcázar, V.; Borrajo, D.; Kissmann, P.; and Edelkamp, S. 2014. SymBA*: A Symbolic Bidirectional A* Planner. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 105–109.

Torralba, Á.; Seipp, J.; and Sievers, S. 2021. Automatic Instance Generation for Classical Planning. In Goldman, R. P.; Biundo, S.; and Katz, M., eds., *Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS 2021)*, 376–384. AAAI Press.

Xie, F.; Müller, M.; and Holte, R. 2014. Jasper: the art of exploration in Greedy Best First Search. In *Eighth International Planning Competition (IPC-8): Planner Abstracts*, 39–42.