Hapori Linear Regression

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

Hapori Linear Regression[1] is a portfolio planner which participated in the optimal, satisficing, and agile tracks of the International Planning Competition (IPC) 2023. It uses linear regression to learn models that predict the best IPC 2018 planner for a given task based on simple planning task features (Ferber and Seipp 2022).

Introduction

No planner excels at all tasks, but each has its individual strength and weaknesses (Roberts and Howe 2009). Thus, planner portfolios try to combine multiple planners to dip into the strength of each one. Delfi (Katz et al. 2018b; Sievers et al. 2019) is an example of an online portfolio. Given a task, Delfi converts it to an image and then uses a convolutional neural network (CNN) to predict the best planner for the task. Delfi is highly successful and won the classical, optimal track of the International Planning Competition (IPC) 2018. Delfi is not only successful, but also unexplainable. Experts understand neither which properties of a planning tasks can be seen in the constructed images nor which rules the CNN learned. The graph convolution based successor of Delfi (Ma et al. 2020) improved upon the first issue, but the second one remained. Ferber and Seipp (2022) successfully train explainable portfolios with a similar performance to Delfi. One of the portfolio selector methods by Ferber and Seipp (2022) is based on linear regression, which receives a set of numeric, explainable features as input and outputs the name of the planner to execute. Here, we retrained this portfolio selector using a larger set of benchmarks tasks and the planners from the IPC 2018.

Method

Let $\mathcal{O}$ be the set of possible observations (in our case the set of all possible PDDL tasks). Let F be a list of *numeric features* such that each feature $f \in F$ is a function $f : \mathcal{O} \rightarrow \mathbb{R}$. A feature vector $\vec{x} \in \mathbb{R}^{|F|}$ for an observation o holds the evaluation of the features on the observation, i.e. $\vec{x}_i = F_i(o)$ for $1 \leq i \leq |F|$. A *linear regression* (Galton 1886) model m predicts a *single* continous value for a feature vector. It is a function $m : \mathbb{R}^{|F|} \rightarrow \mathbb{R}$ of the form $m(\vec{x}) = \vec{x} \cdot \vec{w} + b = [\vec{x}\ 1] \cdot \vec{w}'$, where $\vec{w} \in \mathbb{R}^{|F|}$ ($\vec{w}' \in \mathbb{R}^{|F|+1}$) i.e. it produces a linear combination of the feature values plus a bias term. To train a linear regression model, we require training data, a list of feature vector - outcome pairs $[\langle \vec{x}_i, y_i \rangle]_{i=1..N}$, and a loss function $l : \mathbb{R}^2 \rightarrow \mathbb{R}$. Then, we try to find the weights $\vec{w}'$ with the least loss, i.e.

$$\underset{\vec{w}' \in \mathbb{R}^{|F|+1}}{\arg\min} = \sum_{i=1}^{N} l(m(\vec{x}_i), y_i). \tag{1}$$

Like Ferber and Seipp (2022), we use the mean squared error as loss function. This allows us to compute an analytical solution to the problem:

$$\vec{w}' = (X^T X)^{-1} X^T \vec{y}, \text{with} \tag{2}$$

$$X = [\vec{x}_1^T \ \dots \ \vec{x}_N^T]^T \tag{3}$$

$$\vec{y} = [y_1 \ \dots \ y_N]^T \tag{4}$$

Let T be a list of N tasks and P be a set of planners. For every task $t = T_i$, we compute the feature vector $\vec{x}_i$. For every tasks $t = T_i$ and every planner $p \in P$, we set $y_{i,p}$ to 1, if the planner p solves the task t within the resource limits and to 0 otherwise. Now we train for every planner a model m_p which estimates the likelyhood that the planner p solves a task. We use the matrix X as defined in Equation 3 and $\vec{y}_p = [y_{i,p}]_{i=1..N}$.

Given a new task t', our portfolio executes the planner p with the highest estimated likelyhood of solving the task t, i.e.

$$\underset{p \in P}{\arg\max}\, m_p(\vec{x}_{t'}). \tag{5}$$

Components and Training Data

Planners. As the pool of planners for our portfolios to choose from, we used all planners from the IPC 2018. If an IPC 2018 planner was already a portfolio, we used its component planners instead. We only considered each planner once (some portfolios included planners that were also submitted separately and several portfolios included the same planners).

[1]Hapori is the Maori word for community.

For the optimal track, we had to exclude maplan-1, maplan-2, and MSP because they use CPLEX, and Complementary1 because it generates suboptimal solutions. Furthermore, the FDMS planners and Metis1 were covered by Delfi 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)

All planners participating in the 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 had to exclude alien 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 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)
- LAPKT-DUAL-BFWS, LAPKT-POLYNOMIAL-BFWS, LAPKT-DFS+, and LAPKT-BFWS-Preference (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 Runtime. For training the portfolios, we used all tasks and domains from previous IPCs, from Delfi (Katz et al. 2018b), and from the 21.11 Autoscale collection Torralba, Seipp, and Sievers (2021), leading to a set of 92 domains with 7330 tasks. We used Downward Lab (Seipp et al. 2017) to run all planners on 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 stored its outcome (plan found, out of memory, out of time, task not supported by planner, 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 online available.[2] As training data for our optimal (respectively satisficing/agile) portfolios, we selected 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) remaining tasks.

Features. Ferber and Seipp (2022) showed that their models performed best when trained on the 49 PDDL features of Fawcett et al. (2014). Thus, we also train our models on those features. Among others, those include the number of objects, the number of actions, and the mean number of parameters per predicate. For each task in our benchmark collection, the PDDL features are also available online.

Executing Predictive Portfolios

Given a task, the portfolio selector computes the values of its input features. Then, it evaluates the output of the trained model with respect to the values of the features. Next, it interprets the model output, e.g., if the model directly predicts a planner, then this planner is selected; if it predicts for each planner the probability that it solves the given task, then the planner with highest probability is selected. Finally, it executes the the selected planner for the whole time limit.

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.

References

Fawcett, C.; Vallati, M.; Hutter, F.; Hoffmann, J.; Hoos, H.; and Leyton-Brown, K. 2014. Improved Features for Runtime Prediction of Domain-Independent Planners. In Chien, S.; Fern, A.; Ruml, W.; and Do, M., eds., *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling (ICAPS 2014)*, 355–359. AAAI Press.

Ferber, P.; and Seipp, J. 2022. Explainable Planner Selection for Classical Planning. In Honavar, V.; and Spaan, M., eds., *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022)*, 9741–9749. AAAI Press.

[2]URL to be published

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.

Galton, F. 1886. Regression Towards Mediocrity in Hereditary Stature. *Journal of the Anthropological Institute of Great Britain and Ireland*, 15: 246–263.

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.

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.

Ma, T.; Ferber, P.; Huo, S.; Chen, J.; and Katz, M. 2020. Online Planner Selection with Graph Neural Networks and Adaptive Scheduling. In Conitzer, V.; and Sha, F., eds., *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*, 5077–5084. AAAI Press.

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.

Roberts, M.; and Howe, A. E. 2009. Learning from planner performance. *Artificial Intelligence*, 173: 536–561.

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. 2019. 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.

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.