Structured Convolutional Kernel Networks for Airline Crew Scheduling

Yassine Yaakoubi[1,2,*]   François Soumis[1]   Simon Lacoste-Julien[2,3,†]

[1]GERAD, Polytechnique Montréal   [2]Mila   [3]Université de Montréal

[*]Currently at McGill University   [†]Canada CIFAR AI Chair

ICML 2021

Motivation (Crew Pairing Problem)

Crew pairing problem: Searching for a minimum-cost set of anonymous feasible pairings from the scheduled flights, such that:

- Flights are covered exactly once
- Airline regulations and collective agreements are respected

Existing solution: Commercial-GENCOL-DCA based on column generation and dynamic constraint aggregation by [Yaakoubi et al., 2019, 2020]

Limitation: Cannot produce initial solutions that are sufficiently close to being feasible

Proposed approach: Struct-CKN

Main idea: Impose constraints on the output to propose initial solutions that are feasible enough

Dataset: The flight-connection dataset by Yaakoubi et al. [2019] uses a similarity-based input favoring the use of convolutional architectures

Method: Combine convolutional kernel networks [Mairal et al., 2014, Mairal, 2016, Bietti and Mairal, 2019] in a structured prediction framework that supports constraints on the outputs

Proposed approach: Struct-CKN

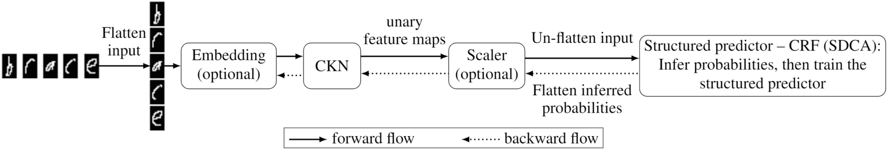

Numerical Results: OCR

Input data: Handwritten words (pre-segmented into characters)
Task: To classify the images into one of the characters

	Test error (%)
SDCA - linear features [Le Priol et al., 2018]	12.0
LSTM [Greff et al., 2016]	4.6
CNN-CRF [Chu et al., 2016]	4.5
SCRBM [Tran et al., 2020]	4.0
NLStruct [Graber et al., 2018]	3.6
Struct-CKN	**3.4**

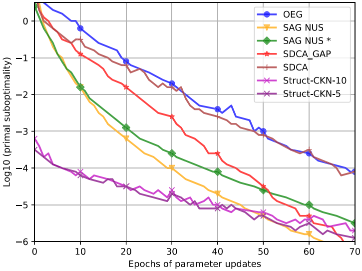

Numerical Results: Crew Pairing

Flight-connection dataset

- Training set: Six monthly crew pairing solutions (50,000 flights/month)
- Test set: A separate test set

CRF graph

- Nodes: Flights
- Arcs: Feasibility of two flights being successive
- True label: Rank of the next flight in the set of possible successors
- Local constraints: Each flight has to be preceded by at most one flight

Numerical Results: Crew Pairing

	Test error (%)	# parameters
CNN [Yaakoubi et al., 2020]	0.32	459 542
CNN-CRF [Chu et al., 2016]	0.38	547 542
Struct-CKN	**0.28**	**15 200**

	#pairings	Cost ($\times 10^8$)	% infeasible pairings
Standard - initial	6 525	37.93	50.56
Standard - feasible	3 226	29.75	
CNN - initial	4 883	24.13	21.05
CNN - feasible	3 855	16.58	
CNN-CRF - initial	4 567	21.15	12.12
CNN-CRF - feasible	4 010	16.15	
Struct-CKN - initial	4 515	20.29	**11.07**
Struct-CKN - feasible	**4 015**	16.46	

Numerical Results: Crew Pairing

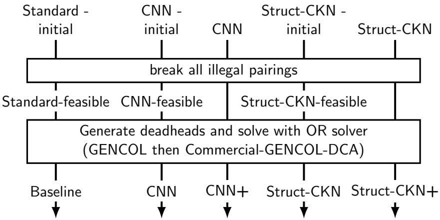

	Solution cost ($\times 10^6$ **\$**)	Cost of global constraints ($\times 10^5$)	Number of deadheads	Total time (hours)
Baseline [Desaulniers et al., 2020]	20.64	21.27	992	45.92
CNN [Yaakoubi et al., 2020]	18.88	4.66	1014	95.72
Struct-CKN	**18.68**	**4.20**	**915**	64.48
CNN+	18.62	3.34	997	126.62
Struct-CKN+	**17.15**	**0.59**	**583**	41.44

Conclusions and future work

Take-aways

- Struct-CKN outperforms state-of-the-art methods on primal sub-optimality and test accuracy on the OCR dataset
- Using Struct-CKN to wamr-start the solver reduces the solution cost by 17% (a gain of millions of dollars) and the cost of global constraints by 97%

Future work

- Combine deep structured methods with various operations research methods
- Design new reactive/learning metaheuristics that learn to guide the search for better solutions in real-time

Structured Convolutional Kernel Networks for Airline Crew Scheduling

Yassine Yaakoubi[1,2*] François Soumis[1] Simon Lacoste-Julien[2,3†]

[1]GERAD, Polytechnique Montréal [2]Mila [3]Université de Montréal

[*]Currently at McGill University [†]Canada CIFAR AI Chair

ICML 2021

Bibliography I

Alberto Bietti and Julien Mairal. Group invariance, stability to deformations, and complexity of deep convolutional representations. *JMLR*, 2019.

Xiao Chu, Wanli Ouyang, hongsheng Li, and Xiaogang Wang. Crf-cnn: Modeling structured information in human pose estimation. In *Advances in Neural Information Processing Systems 29*, pages 316–324. Curran Associates, Inc., 2016.

Guy Desaulniers, François Lessard, Saddoune Mohammed, and Soumis François. Dynamic constraint aggregation for solving very large-scale airline crew pairing problems. *Les Cahiers du GERAD*, G–2020(21), 2020.

Colin Graber, Ofer Meshi, and Alexander Schwing. Deep structured prediction with nonlinear output transformations. In *Advances in Neural Information Processing Systems*, 2018.

Bibliography II

Klaus Greff, Rupesh K Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 28(10), 2016.

Rémi Le Priol, Alexandre Piché, and Simon Lacoste-Julien. Adaptive stochastic dual coordinate ascent for conditional random fields. *Conference on Uncertainty in Artificial Intelligence*, 2018.

J. Mairal, P. Koniusz, Z. Harchaoui, and C. Schmid. Convolutional kernel networks. *Advances in Neural Information Processing Systems*, 2014.

Julien Mairal. End-to-end kernel learning with supervised convolutional kernel networks. In *Advances in Neural Information Processing Systems*, 2016.

Son N Tran, Artur d'Avila Garcez, Tillman Weyde, Jie Yin, Qing Zhang, and Mohan Karunanithi. Sequence classification restricted boltzmann machines with gated units. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

Bibliography III

Yassine Yaakoubi, François Soumis, and Simon Lacoste-Julien. Flight-connection prediction for airline crew scheduling to construct initial clusters for OR optimizer. *Les Cahiers du GERAD*, G–2019(26), 2019.

Yassine Yaakoubi, François Soumis, and Simon Lacoste-Julien. Machine learning in airline crew pairing to construct initial clusters for dynamic constraint aggregation. *EURO Journal on Transportation and Logistics*, 2020. ISSN 2192-4376. doi: 10.1016/j.ejtl.2020.100020.