

---

# Structured Convolutional Kernel Networks for Airline Crew Scheduling

---

Yassine Yaakoubi<sup>1 2</sup> François Soumis<sup>1</sup> Simon Lacoste-Julien<sup>† 2 3</sup>

## Abstract

Motivated by the needs from an airline crew scheduling application, we introduce structured convolutional kernel networks (Struct-CKN), which combine CKNs from Mairal et al. (2014) in a structured prediction framework that supports *constraints* on the outputs. CKNs are a particular kind of convolutional neural networks that approximate a kernel feature map on training data, thus combining properties of deep learning with the non-parametric flexibility of kernel methods. Extending CKNs to structured outputs allows us to obtain useful initial solutions on a flight-connection dataset that can be further refined by an airline crew scheduling solver. More specifically, we use a flight-based network modeled as a general conditional random field capable of incorporating local constraints in the learning process. Our experiments demonstrate that this approach yields significant improvements for the large-scale crew pairing problem (50,000 flights per month) over standard approaches, reducing the solution cost by 17% (a gain of millions of dollars) and the cost of global constraints by 97%.

## 1. Introduction

Since crew costs are the second-highest spending source for air passenger carriers, crew scheduling is of crucial importance for airlines. The crew pairing problem (CPP) searches for a minimum-cost set of anonymous feasible pairings (rotations) from the scheduled flights, such that all flights are covered exactly once, and all airline regulations and collective agreements are respected. The complexity of this problem lies in the large number of possible pairings, as the selection of pairings at minimal cost—a large integer programming problem—cannot be performed using standard solvers. Seeking to obtain an efficient algorithm for

large-scale monthly CPPs (up to 50,000 flights) and building on the column generation-based solver by Desaulniers et al. (2020), Yaakoubi et al. (2020) proposed Commercial-GENCOL-DCA, an improved solver starting with an aggregation, in clusters, of flights. The initial aggregation partition permits replacing all flight-covering constraints of flights in a cluster by a single constraint, thus allowing the solver to cope with larger instances. Initial clusters can either be extracted from the initial solution (Desaulniers et al., 2020) or given separately, as in Yaakoubi et al. (2020), where authors used convolutional neural networks (CNN) to solve the flight-connection problem (a supervised multi-class classification problem). The objective of this problem is to predict the next flight that a crew follows in its schedule given the previous flight. They used CNN to harness the spatial locality (localized spatial features) and used a similarity-based input, where neighboring factors have similar features. By passing initial clusters of flights to the CPP solver, the reported reduction of solution cost averages between 6.8% and 8.52%, mainly due to the reduction in the cost of global constraints between 69.79% and 78.11%. The cost of global constraints refers to the penalties incurred when the workload is not fairly distributed among the bases in proportion to the available personnel at each base.

However, a major weakness of their approach is that they can only produce initial clusters and not an initial solution, since they use a greedy predictor making one prediction at a time (predicting sequentially the next flight given only the previous flight). This prevents the predictor from incorporating constraints on the output and the produced solutions cannot be used as initial solutions for the solver as they are not sufficiently close to being feasible. A pairing is deemed feasible if it satisfies safety rules and collective agreement rules (Kasirzadeh et al., 2017); examples include minimum connection time between two flights, minimum rest time, and maximum number of duties in a pairing. By providing an initial solution, we not only accelerate the optimization process and calculate the feasibility of proposed pairings, but we also propose clusters similar to the initial solution, thus reducing the degree of incompatibility between current solution and proposed pairings (Elhallaoui et al., 2010).

In this paper, we address this lack of constraint modeling, while still enabling the use of a convolutional architecture. For this purpose, we investigate the convolutional kernel

---

<sup>1</sup>GERAD, Polytechnique Montréal, Canada <sup>2</sup>Mila, Canada <sup>3</sup>Department of Computer Science and Operations Research, Université de Montréal, Canada † Canada CIFAR AI Chair. Correspondence to: Yassine Yaakoubi <YassineYaakoubi@outlook.com>.

network (CKN), an approximation scheme similar to CNN proposed by Mairal et al. (2014). To bypass the major limitations in Yaakoubi et al. (2020), we incorporate local constraints on the outputs (imposing that each flight has to be preceded by at most one flight). We harness the spatiotemporal structure of the CPP problem by combining kernel methods and structured prediction. The outputs start the optimizer and solve a large-scale CPP, where small savings of a mere 1% translate into an increase of annual revenue for a large airline by dozens of millions of dollars. Note that, to the best of our knowledge, we are not aware of any ML approach that can directly solve the CPP (which has complex airline-dependent costs and constraints that are not necessarily available to the ML system at train time).

We thus consider instead to use the ML system to propose good initial clusters and an initial solution for the CPP solver. The results of training on the flight-connection dataset (Yaakoubi et al., 2019), a flight-based network structure modeled as a general conditional random field (CRF) graph, demonstrate that the proposed predictor is more suitable than other methods. Specifically, it is more stable than CNN-based predictors and extensive tuning is not required, in that no Bayesian optimization (to find a suitable configuration) is needed, as we observe in our experiments. This is crucial to integrate ML into a solver for the CPP or any real-world scheduling problem. Note that unlike recurrent neural networks (RNNs) or neural networks by (Yaakoubi et al., 2020) which cannot produce initial solutions that are sufficiently close to being feasible, the proposed predictor incorporates local constraints in the learning process. Furthermore, note that while previous studies focused on using ML (either through imitation learning or reinforcement learning) to solve small-scale CO problems such as vehicle routing ( $\leq 100$  customers) and airline crew scheduling ( $\leq 714$  flights) problems, we use the proposed predictor to warm-start a monthly CPP solver (up to 50,000 flights). For an extensive literature review on using machine learning for combinatorial optimization, see Bengio et al. (2020).

**Contributions.** Bridging the gap between kernel methods and neural networks, we propose the structured convolutional kernel network (Struct-CKN)<sup>1</sup>. We first sanity check the approach on the OCR dataset (Taskar et al., 2004) yielding a test accuracy comparable to state of the art. Then, to warm-start an airline crew scheduling solver, we apply the proposed method on a flight-connection dataset, modeled as a general CRF capable of incorporating local constraints in the learning process. We show that the constructed solution outperforms other approaches in terms of test error and feasibility, an important metric to initialize the solver. The predicted solution is fed to the solver as an initial solu-

tion and initial clusters, to solve a large-scale CPP (50,000 flights). Our experiments demonstrate that this approach yields significant improvements, reducing the solution cost by 17% (a gain of millions of dollars) and the cost of global constraints by 97%, compared to baselines.

**Outline.** The remainder of this paper is structured as follows. Section 2 describes related methods. Section 3 presents CKNs. CRFs are outlined in Section 4. Section 5 presents Struct-CKN. Section 6 reports Computational results on OCR dataset, flight-connection dataset, and CPP.

## 2. Related Work

Upon a succinct review of previous work to compare available approaches in the literature to Struct-CKN, we argue for the use of the latter on the flight-connection dataset to solve CPPs.

Combining networks and energy-based models is a well-known approach since the 1990s. For instance, Bottou (2012) introduced graph transformer networks trained end-to-end using weighted acyclic directed graphs to represent a sequence of digits in handwritten character recognition. Furthermore, inspired by Q-learning, Gygli et al. (2017) used an oracle value function as the objective for energy-based deep networks, and Belanger & McCallum (2016) introduced structured prediction energy networks (SPENs) to address the inductive bias and to learn discriminative features of the structured output automatically. By assigning a score to an entire prediction, SPENs take into consideration high-order interactions between predictors using minimal structural assumptions. Nevertheless, due to the non-convexity, optimizing remains challenging, which may cause the learning model to get stuck in local optima. Another approach is to move step by step and predict one output variable at a time by applying the information gathered from previous steps. The linking between the steps is learned using a predefined order of input variable where the conditional is modeled with RNNs (Zheng et al., 2015b). Although this method has achieved impressive results in machine translation (Leblond et al., 2017), its success ultimately depends on the neural network’s ability to model the conditional distribution and it is often sensitive to the order in which input data is processed, particularly in large-size graphs, as in CPPs (50,000 nodes).

In contrast to these approaches, instead of using continuous relaxation of output space variables (Belanger & McCallum, 2016), Struct-CKN uses supervised end-to-end learning of CKNs and CRF-based models. Accordingly, any of the existing inference mechanisms—from belief propagation to LP relaxations—can be applied. This allows us to naturally handle general problems that go beyond multi-label classification, and to apply standard structured loss functions (instead of extending them to continuous variables,

<sup>1</sup>The code is available at the following link: <https://github.com/Yaakoubi/Struct-CKN>

as in the case of SPENs). More importantly, Struct-CKN allows us to apply our method to a large-scale CRF graph containing up to 50,000 nodes. Furthermore, in contrast to methods in the literature (e.g., CNN-CRF (Chu et al., 2016), CRF-RNN (Zheng et al., 2015a), and deep structured models (Chen et al., 2015)), it has far fewer parameters, thus bypassing the need for extensive tuning.

Finally, note that in recent papers, convolutional graph neural networks (ConvGNNs) (Kipf & Welling, 2016) are used either (1) to warm-start a solver (trained under the imitation learning framework) (Owerko et al., 2020), (2) to solve the optimization problem end-to-end (Khalil et al., 2017), or (3) to guide an optimization process (variable selection in branch-and-bound Gasse et al. (2019)), and ConvGNNs might appear to be a good candidate-solution for CPPs when coupled with a CRF layer to impose constraints on the output. However, in the case of graphs with up to 50,000 nodes, the number of parameters used by ConvGNN and the computational limitations prevents us from considering it. In fact, we are not aware of any prior work where ConvGNNs were used at this scale. Furthermore, the implicit motivation for our proposed approach is an end-to-end solution method that can (1) harness the predictive capabilities of the ML predictor and the decomposition capacity of the solver (Yaakoubi et al., 2020), and (2) can be used on a standard machine with no specific resource requirements, to replace existing solvers in the industry. Future research will look into the possibility of integrating a distributed version of ConvGNNs into the proposed framework.

### 3. Convolutional Kernel Networks

CKN is a particular type of CNN that differs from the latter in the cost function to be optimized to learn filters and in the choice of non-linearities. We review CKNs, with the same notation as in Mairal (2016); Bietti & Mairal (2019). For further detail, see Appendix A.

#### 3.1. Unsupervised Convolutional Kernel Networks (Mairal et al., 2014)

We consider an image  $I_0 : \Omega_0 \rightarrow \mathbb{R}^{p_0}$ , where  $p_0$  is the number of channels, e.g.,  $p_0 = 3$  for RGB, and  $\Omega_0 \subset [0, 1]^2$  is a discrete set of pixel locations. Given two image patches  $x, x'$  of size  $e_0 \times e_0$ , represented as vectors in  $\mathbb{R}^{p_0 e_0^2}$ , we define a kernel  $K_1(x, x') = \|x\| \|x'\| \cdot \kappa_1(\langle \frac{x}{\|x\|}, \frac{x'}{\|x'\|} \rangle)$  if  $x, x' \neq 0$  and 0 otherwise, where  $\|\cdot\|$  and  $\langle \cdot, \cdot \rangle$  denote the Euclidian norm and inner-product, respectively, and  $\kappa_1(\langle \cdot, \cdot \rangle)$  is a dot-product kernel on the sphere. We have implicitly defined the reproducing kernel Hilbert space (RKHS)  $\mathcal{H}_1$  associated to  $K_1$  and a mapping  $\varphi_1 : \mathbb{R}^{p_0 e_0^2} \rightarrow \mathcal{H}_1$ .

First, we build a database of  $n$  patches  $\mathbf{x}_1, \dots, \mathbf{x}_n$  randomly extracted from various images and normalized to

have unit  $\ell_2$ -norm. Then, we perform a spherical  $K$ -means algorithm (Buchta et al., 2012), acting as a Nyström approximation, to obtain  $p_1$  centroids  $\mathbf{z}_1, \dots, \mathbf{z}_{p_1}$  with unit  $\ell_2$ -norm. Given a patch  $\mathbf{x}$  of  $I_0$ , the projection of  $\varphi_1(\mathbf{x})$  onto  $\mathcal{F}_1 := \text{Span}(\varphi_1(\mathbf{z}_1), \dots, \varphi_1(\mathbf{z}_{p_1}))$  admits a natural parametrization given in (1) where  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_{p_1}]$ , and  $\kappa_1$  is applied pointwise to its arguments.

$$\Gamma_1(\mathbf{x}) := \|\mathbf{x}\| \kappa_1(\mathbf{Z}^\top \mathbf{Z})^{-1/2} \kappa_1\left(\mathbf{Z}^\top \frac{\mathbf{x}}{\|\mathbf{x}\|}\right) \quad (1)$$

if  $\mathbf{x} \neq 0$  and 0 o.w.

Consider all overlapping patches of  $I_0$ . We set  $M_1(z) = \Gamma_1(\mathbf{x}_z)$ ,  $z \in \Omega_0$  where  $\mathbf{x}_z$  is the patch from  $I_0$  centered at pixel location  $z$ . The spatial map  $M_1 : \Omega_0 \rightarrow \mathbb{R}^{p_1}$  thus computes the quantities  $\mathbf{Z}^\top \mathbf{x}$  for all patches  $\mathbf{x}$  of image  $I$  (spatial convolution after mirroring the filters  $\mathbf{z}_j$ ), then applies the pointwise non-linear function  $\kappa_1$ .

The previous steps transform the image  $I_0 : \Omega_0 \rightarrow \mathbb{R}^{p_0}$  into a map  $M_1 : \Omega_0 \rightarrow \mathbb{R}^{p_1}$ . Then, the CKNs involve a pooling step to gain invariance to small shifts, leading to another finite-dimensional map  $I_1 : \Omega_1 \rightarrow \mathbb{R}^{p_1}$  with a smaller resolution:  $I_1(z) = \sum_{z' \in \Omega_0} M_1(z') e^{-\beta_1 \|z' - z\|_2^2}$ ,  $z \in \Omega_1$ , where  $\beta_1$  is a subsampling factor. We build a multilayer image representation by stacking and composing kernels. Similarly to the first CKN layer transforming  $I_0 : \Omega_0 \rightarrow \mathbb{R}^{p_0}$  to the map  $I_1 : \Omega_1 \rightarrow \mathbb{R}^{p_1}$ , we apply the same procedure to obtain  $I_2 : \Omega_2 \rightarrow \mathbb{R}^{p_2}$ , where  $p_2$  is the number of centroids in the second layer, then  $I_3 : \Omega_3 \rightarrow \mathbb{R}^{p_3}$ , etc.

#### 3.2. Supervised Convolutional Kernel Networks (Mairal, 2016)

Let  $I_0^1, I_0^2, \dots, I_0^n$  be the training images with respective labels  $y_1, \dots, y_n$  in  $\{-1; +1\}$  for binary classification. We also have  $L : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ , a convex smooth loss function. Given a positive definite kernel  $K$  on images, the classical empirical risk minimization formulation consists of finding a prediction function in the RKHS  $\mathcal{H}$  associated to  $K$  by minimizing the objective  $\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n L(y_i, f(I_0^i)) + \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2$ , where the parameter  $\lambda$  controls the smoothness of the prediction function  $f$  with respect to the geometry induced by the kernel, hence regularizing and reducing overfitting.

After training a CKN with  $k$  layers, such a positive definite kernel  $K_{\mathcal{Z}}$  may be defined as in (2) where  $I_k, I'_k$  are the  $k$ -th finite-dimensional feature maps of  $I_0$  and  $I'_0$ , respectively, and  $f_k, f'_k$  are the corresponding maps in  $\Omega_k \rightarrow \mathcal{H}_k$ , which have been defined in Section 3.1.

$$K_{\mathcal{Z}}(I_0, I'_0) = \sum_{z \in \Omega_k} \langle f_k(z), f'_k(z) \rangle_{\mathcal{H}_k} = \sum_{z \in \Omega_k} \langle I_k(z), I'_k(z) \rangle \quad (2)$$

The kernel  $K_{\mathcal{Z}}$  is also indexed by  $\mathcal{Z}$ , representing network parameters (subspaces  $\mathcal{F}_1, \dots, \mathcal{F}_k$ , or equivalently the set of filters  $Z_1, \dots, Z_k$ ). Then, the formulation becomes as in (3) where  $\|\cdot\|_F$  is the Frobenius norm extending the Euclidean norm to matrices and, with an abuse of notation, the maps  $I_k^i$  are seen as matrices in  $\mathbb{R}^{p_k \times |\Omega_k|}$ . Then, the supervised CKN formulation consists of jointly minimizing (3) w.r.t.  $W$  in  $\mathbb{R}^{p_k \times |\Omega_k|}$  and with respect to the set of filters  $Z_1, \dots, Z_k$ , whose columns are constrained to be on the Euclidean sphere.

$$\min_{W \in \mathbb{R}^{p_k \times |\Omega_k|}} \frac{1}{n} \sum_{i=1}^n L(y_i, \langle W, I_k^i \rangle) + \frac{\lambda}{2} \|W\|_F^2 \quad (3)$$

## 4. Graph-Based Learning

In structured prediction, models are typically estimated with surrogate structured loss minimization, such as with structured SVM (SSVM) or CRFs. We used CRFs for the structured prediction that we briefly review below. We also tested SSVM integration instead of CRFs, but it yielded slightly worse results, see Appendices B.2 and C for details and experimental results using SSVM.

A CRF models the conditional probability of a structured output  $y \in \mathcal{Y}$  given an input  $x \in \mathcal{X}$  where the probability to observe  $y$  when  $x$  is observed is  $p(y|x; w) \propto \exp(\langle w, F(x, y) \rangle)$ ,  $F$  is the feature mapping, and  $w$  is the vector of weights (to be learned). The CRF predictor of  $y$  when  $x$  is observed is:  $h_w(x) = \arg \max_{y \in \mathcal{Y}} \langle w, F(x, y) \rangle$ . The CRF primal problem formulation is shown in (4), where  $\mathcal{L}^{CRF}$  denotes the negative log likelihood loss.

$$\min_w \lambda \|\omega\|^2 + \frac{1}{n} \sum_{i=1}^N \mathcal{L}^{CRF}(x_i, y_i; \omega) \quad (4)$$

To optimize CRFs, we use stochastic dual coordinate ascent (SDCA) (Shalev-Shwartz & Zhang, 2013; 2014) as Le Priol et al. (2018) showed it yielded state-of-the-art results for CRFs. Although one can also use the stochastic average gradient (SAG) algorithm (Schmidt et al., 2015) or the online exponentiated gradient (OEG) algorithm, an advantage of SDCA over OEG (and SAG) is that it enables performing an “exact” line search with only one call to the marginalization oracle. We now rewrite (4) using the notation for the SDCA setup for multi-class classification (Shalev-Shwartz & Zhang, 2014). Denote  $M_i = |\mathcal{Y}_i|$  the number of labelings for sequence  $i$ . Denote  $A_i$  the matrix whose columns are the corrected features  $\{\psi_i(y) := F(x_i, y_i) - F(x_i, y)\}_{y \in \mathcal{Y}_i}$ . Denote also  $\Phi_i(s) := \log(\sum_{y \in \mathcal{Y}_i} \exp(s_y))$  the log-partition function for the scores  $s \in \mathbb{R}^{M_i}$ . Since the negative log-likelihood can be written as  $-\log(p(y_i|x_i; w)) = \Phi_i(-A_i^\top w)$  (Mur-

phy, 2012), the primal objective function to minimize over  $w \in \mathbb{R}^d$  becomes  $P(w) := \frac{\lambda}{2} \|w\|_2^2 + \frac{1}{n} \sum_{i=1}^n \Phi_i(-A_i^\top w)$ . This minimization problem has an equivalent Fenchel convex dual problem. Denote  $\Delta_M$  the probability simplex over  $M$  elements. Denote  $\alpha_i \in \Delta_{M_i}$  the set of dual variables for a given  $x_i$ , we define the conjugate weight function  $\hat{w}$  as follows:  $\hat{w}(\alpha) = \frac{1}{\lambda n} \sum_i A_i \alpha_i = \frac{1}{\lambda n} \sum_{i=1}^n \mathbb{E}_{y \sim \alpha_i}(\psi(y))$ . We can show that  $\hat{w}(\alpha^*) = w^*$  where  $w^*$  and  $\alpha^*$  are respectively the optimal primal parameters and the optimal dual parameters. As such, we can also define the primal sub-optimality as:  $P(w) - P(w^*)$ .

Le Priol et al. (2018) adapted SDCA to CRF by considering marginal probabilities over cliques of the graphical model. Because the dual variable  $\alpha_i$  is exponentially large in input size  $x_i$ ,  $\alpha$  is replaced by  $\mu = (\mu_1, \dots, \mu_n)$ , where  $\mu_i \in \Pi_C^{\Delta_C}$  is the concatenation of all the clique marginal vectors for sample  $i$ . Given its state-of-the-art performance, we decided to use it, although we will compare it with other optimizers in Section 6 (see Appendix B.1 for further details).

## 5. The Struct-CKN Framework

As in Figure 1, the Struct-CKN model consists of two components intended to train CRFs: (1) the CKN and (2) the structured predictor, using CRF loss, and SDCA. Upon initializing the CKN layers and the structured predictor, for each iteration, we pass the input image through the CKN multilayers. The last map of CKN is passed through to the structured predictor to infer probabilities, which are employed to train CKN weights by backpropagating using rules in Mairal (2016).

---

### Algorithm 1 Training the Struct-CKN Model

---

- 1: Initialize CKN parameters in unsupervised manner as described in Sec. 3.1
  - 2: Initialize structured predictor and CRF model as in Alg. 2 (see Appendix B.1, steps 1-2)
  - 3: **for**  $t = 0 \dots$  **do**
  - 4:   For each input, construct an unary feature map, as described in Sec. 3.1
  - 5:   (Optional) Center and rescale these representations to have unit  $\ell_2$ -norm on average
  - 6:   Infer probabilities by providing image map as input to structured predictor
  - 7:   Train the structured predictor using the feature map as an input for  $n_{Ep}$  epochs
  - 8:   Use the inferred probabilities to compute the gradient by using the chain rule (backpropagation) and update the CKN weights (Mairal, 2016)
  - 9: **end for**
- 

To do inference for CRF models with a small number of nodes (as in Section 6.1), we use max-product belief prop-

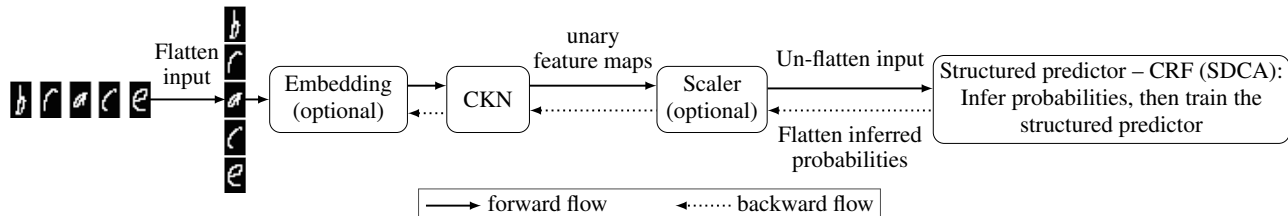


Figure 1. The Architecture Diagram for the Struct-CKN Predictor

agation, since chains can be solved exactly and efficiently. For the flight-connection dataset (as in Section 6.2), since CRF models contain 50,000 nodes, we use AD3 (alternating directions dual decomposition) (Martins et al., 2015) for approximate maximum a posteriori (MAP) inference. It allies the modularity of dual decomposition with the effectiveness of augmented Lagrangian optimization via the alternating directions method of multipliers and has some very interesting features in comparison to other message-passing algorithms. Indeed, AD3 has been empirically shown to reach consensus faster than other algorithms Martins et al. (2015) and outperforms state-of-the-art message-passing algorithms on large-scale problems. Besides, AD3 provides a library of computationally-efficient factors that allow handling declarative constraints within an optimization problem. This is particularly interesting for the CPP use case since we add a constraint to impose that each flight is preceded by one flight at most. To use SDCA (requiring a marginalization oracle) with AD3 (used to do approximate MAP), we propose a simple approximation, using MAP label estimates. See Appendix B.1 for details on integrating SDCA and AD3.

Note that an embedding layer may be used before passing the input through to the CKN layers. Indeed, for categorical variables with a large number of categories, the input matrix is sparse, making the learning process difficult, as the extracted patches have mostly null values. Furthermore, as in Chandra et al. (2017), using deep structured predictors may require using scaling. Specifically, the last map of the “deep” layer needs to be rescaled before being passed to the “structured layer”. We propose to use one of the following scalers within the Scikit-learn library: Min-Max scaler, Normalizer scaler, Standard scaler, and Robust scaler. When using the SDCA optimizer, line search requires computing the entropy of the marginals. Since this is costly, in order to minimize the number of iterations, we used the Newton-Raphson algorithm. This requires storing the logarithm of the dual variable, which may be expensive, so a decent amount of memory should be allocated.

Finally, note that we use a batch-version of the Struct-CKN predictor on the flight-connection dataset (Yaakoubi et al., 2019), where one batch corresponds to one CRF model

Table 1. Test Error on the OCR Dataset

	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	<b>3.4</b>

(CPP instance): (1) We initialize the CKN weights (in an unsupervised manner) and the CRF model sequentially by considering each one of the six instances separately. (2) We “flatten” the input by considering each one of the six instances separately. (3) We center and rescale the representations for all the instances at once. (4) We pass small batches of image maps as inputs to the structured predictor (e.g., 128 image maps) to infer the probabilities and train the structured predictor. Then, we use the inferred probabilities of the batch to update the CKN weights.

## 6. Experiments

In this section, we report the results of experiments using Struct-CKN. First, we sanity-check Struct-CKN on the standard OCR dataset in Section 6.1, showing that it’s comparable to the state of the art. Then, in Section 6.2, we use the proposed predictor on the flight-connection dataset to warm-start Commercial-GENCOL-DCA. We use Pytorch (Paszke et al., 2019) to declare said model and perform operations on a 40-core machine with 384 GB of memory, and use K80 (12 GB) GPUs. The CRF model is implemented using PyStruct (Müller & Behnke, 2014), while the SDCA optimizer is implemented using SDCA4CRF. Scalers are implemented using Scikit-learn (Pedregosa et al., 2011).

### 6.1. OCR - Chain CRF

Each example in the OCR dataset (Taskar et al., 2004) consists of a handwritten word pre-segmented into characters, with each character represented as a  $16 \times 8$  binary image. The task is to classify the image into one of the 26 charac-

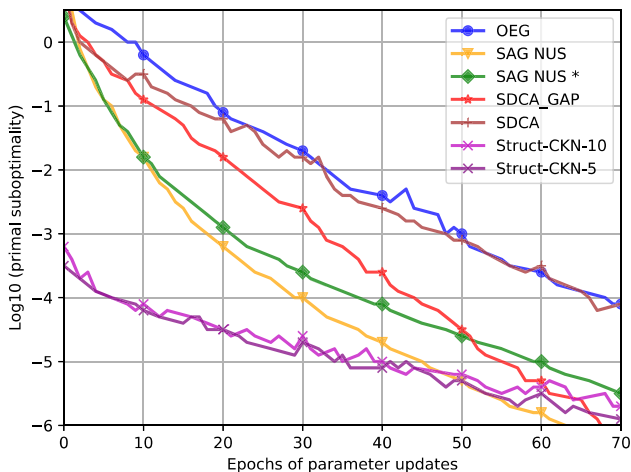


Figure 2. Comparison of Primal Sub-optimality

ters (a–z). It comes with pre-specified folds; one fold is considered the test set, while the rest as the training set, as in max-margin Markov networks (Taskar et al., 2004). Since the CRF optimizers (SAG-NUS, SAG-NUS\* (Schmidt et al., 2015), SDCA, SDCA-GAP (Le Priol et al., 2018), and OEG (Schmidt et al., 2015)) yield similar test errors (11.8–12%), we only report SDCA (with linear features) in Table 1. LSTM (standard two-layer) (Greff et al., 2016) and CNN-CRF (standard two-layer) (Chu et al., 2016) yield comparable results (4.4–4.6%). Sequence Classification Restricted Boltzmann Machine (SCRBM) (Tran et al., 2020) and NLStruct (Graber et al., 2018) provide better results (4.0%, and 3.6%). However, Struct-CKN outperforms all aforementioned methods, reducing test errors to 3.40%.

Note that some structured predictors can lower test error to 1 – 3% (Pérez-Cruz et al., 2007), such as SeaRNN (Leblond et al., 2017), which adapts RNN to the learning-to-search approach. However, such models approximate the cost-to-go for each token by computing the task loss for as many roll-outs as the vocabulary size at each time step, and are thus difficult to scale to real-world datasets (with long sequences or large vocabulary), such as the flight-connection dataset (see Section 6.2). Figure 2 reports primal sub-optimality w.r.t parameter updates (see Appendix C). Struct-CKN outperforms other methods for the first 50 epochs and is comparable to other methods for subsequent epochs. We usually train predictors only for several epochs; thus, the precision of primal sub-optimality (below  $10^{-5}$ ) is negligible.

## 6.2. Airline Crew Scheduling Dataset - Graph CRF

This section reports results of using Struct-CKN to warm-start Commercial-GENCOL-DCA (Yaakoubi et al., 2020) and solve large-scale CPP. Section 6.2.1 presents CPP. Section 6.2.2 outlines the flight-connection prediction problem

and CRF models. Section 6.2.3 reports results of predictions. Section 6.2.4 describes the optimization process and analyzes the feasibility of proposed monthly solutions. Section 6.2.5 reports results of solving CPPs using the solver.

### 6.2.1. CREW PAIRING PROBLEM

The CPP aims to find a set of pairings at minimal cost for each category of the crew and each type of aircraft fleet (Desaulniers et al., 1997). A flight sequence operated by a single crew forms a duty and a consecutive sequence of duty periods is named a pairing. A pairing is deemed feasible if it satisfies safety rules and collective agreement rules (Kasirzadeh et al., 2017), such as:

- minimum connection time between two consecutive flights and minimum rest-time between two duties;
- maximum number of flights per duty and maximum span of a duty;
- maximum number of landings per pairing and maximum flying time in a pairing;
- maximum number of days and maximum number of duties in a pairing.

In addition, CPPs use base constraints (referred to as global constraints) to distribute the workload fairly amongst the bases proportionally to the personnel available at each base. Penalties when the workload is not fairly distributed is called the cost of global constraints. In our approach, the predictor does not consider these very complex airline-dependent constraints as well as the solution cost and the cost of global constraints, as doing so would require more data than are currently available. Whenever a flight appears in more than one pairing, we use deadheads: one crew operates the flight, while the others are transferred between two stations for repositioning.

The CPP has been traditionally modelled as a set partitioning problem, with a covering constraint for each flight and a variable for each feasible pairing (Desaulniers et al., 1997; Kasirzadeh et al., 2017). Formally, we consider  $F$  to be a set of flights that must be operated during a given period and  $\Omega$  to be the set of all feasible pairings that can be used to cover these flights. It is computationally infeasible to list all pairings in  $\Omega$  when solving CPPs with more than hundreds of flights. Therefore, it is not tractable to do so in this context (CPPs with 50,000 flights). For each pairing  $p \in \Omega$ , let  $c_p$  be its cost and  $a_{fp}$ ,  $f \in F$ , be a constant equal to 1 if it contains leg  $f$  and 0 otherwise. Moreover, let  $x_p$  be a binary variable that takes value 1 if pairing  $p$  is selected, and 0 otherwise. Using a set-partitioning formulation, the CPP can be modelled as follows:

$$\underset{x}{\text{minimize}} \quad \sum_{p \in \Omega} c_p x_p \quad (5)$$

$$\text{subject to} \quad \sum_{p \in \Omega} a_{fp} x_p = 1 \quad \forall f \in F \quad (6)$$

$$x_p \in \{0, 1\} \quad \forall p \in \Omega \quad (7)$$

The objective function (5) minimizes the total pairing costs. Constraints (6) ensure each leg is covered exactly once, and constraints (7) enforce binary requirements on the pairing variables. The methodology to solve CPPs depends on the size of the airline’s network, rules, collective agreements, and cost structure (Yen & Birge, 2006). Since the 1990s, the most prevalent method has been column generation inserted in branch-&-bound (Desaulniers et al., 1997). This algorithm was combined with multiple methods in Desaulniers et al. (2020) to solve large-scale CPPs. Yaakoubi et al. (2020) proposed Commercial-GENCOL-DCA with a dynamic control strategy and used CNNs to develop initial monthly crew pairings. Because the constructed solution contained too many infeasible pairings, it was passed to the solver only as initial clusters and a generic standard initial solution was used. In this work, we first use Struct-CKN to construct initial monthly crew pairings passed to the solver both as initial clusters and an initial solution. First, we use the GENCOL solver (used to assign crews to under-covered flights),<sup>2</sup> then we run Commercial-GENCOL-DCA. Because solvers can only handle a few thousand flights, the windowing approach is used: the month is divided into multiple windows, where each window is solved (sequentially) while flights in pairings from previous windows are frozen. GENCOL requires using two-day windows and one-day overlap period, while Commercial-GENCOL-DCA permits to use one-week windows and two-day overlap period. The latter starts with an aggregation, in clusters, of flights. The initial aggregation partition permits replacing all flight-covering constraints of the flights in a cluster by a single constraint, allowing to cope with larger instances.

### 6.2.2. PREDICTION PROBLEM FORMULATION

We aim to provide the CPP solver with an initial solution and initial clusters. Using the flight-connection dataset built in Yaakoubi et al. (2019), the flight-connection prediction problem is a multi-class classification problem, formulated as follows: “Given the information about an incoming flight in a specific connecting city, choose among all possible departing flights from this city the one that the crew should follow” (see Appendix D). Thus, each input contains information on the previous flight performed by the crew as well as information on all possible next flights (up to 20 candi-

dates), sorted by departure time, and the output (class) is the rank of the flight performed by the crew in past solutions. In Yaakoubi et al. (2019), the authors propose a similarity-based input where neighboring factors have similar features, allowing the use of CNNs. We extend their approach with our Struct-CKN model, thus starting the CPP solver with an initial solution, and not only the initial clusters.

The training set in the flight-connection dataset consists of six monthly crew pairing solutions (50,000 flights per month) and the test set is a benchmark that airlines use to decide on the commercial solver to use. Each flight is characterized by the cities of origin and destination, the aircraft type, the flight duration, and the departure and arrival time. For each incoming flight, the embedded representation of the candidate next flights is concatenated to construct a similarity-based input, where neighboring factors have similar features. The intuition is that each next flight is considered a different time step, enabling the use of convolutional architecture across time. Yaakoubi et al. (2019) compare multiple predictors on the flight-connection dataset and empirically confirm this intuition (see Appendix D).

The dataset is used to define a pairwise CRF on a general graph where each example consists of a flight-based network structure with nodes corresponding to flights and *arcs representing the feasibility of two flights being successive*. Each flight corresponds to a node connected to nodes that are possible successors/predecessors; the true label is the rank of the next flight in the set of sorted possible successors. We impose local constraints to the output by imposing that each flight has to be preceded by at most one flight (using a XOR constraint, which contributes  $-\infty$  to the potential). Unlike for the OCR dataset where “max-product” is used for belief propagation, in this section, we use AD3 (alternating directions dual decomposition) (Martins et al., 2015) for approximate maximum a posteriori (MAP) inference.

### 6.2.3. RESULTS ON THE FLIGHT-CONNECTION DATASET

As Struct-CKN has fewer hyperparameters than CNNs, we observed that Struct-CKN is more stable than CNNs in our experiments, in that it does not require Bayesian optimization to find a good architecture, thus justifying our use of CKN (see Appendix C). Table 2 reports the number of parameters and the test error on the flight-connection dataset using (1) CNNs and Bayesian optimization to search for the best configuration of hyperparameters (Yaakoubi et al., 2020); (2) standard CNN-CRF (with non-exhaustive hyperparameter tuning) (Chu et al., 2016); and (3) Struct-CKN. Struct-CKN outperforms both CNN and CNN-CRF while having far fewer parameters (97% fewer parameters). While Bayesian optimization can be used to fine-tune hyperparameters, this is not feasible in a real-case usage scenario, as practitioners cannot perform it each time new data become

<sup>2</sup><http://www.ad-opt.com/optimization/why-optimization/column-generation/>

Table 2. Test Error on the Flight-connection Dataset (Yaakoubi et al., 2019)

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

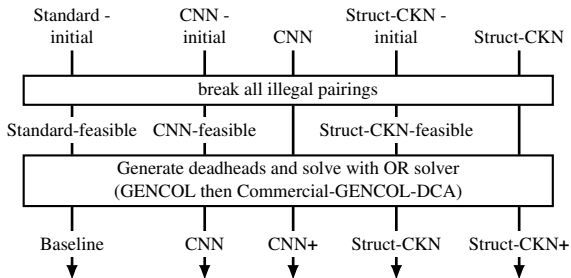


Figure 3. The Optimization Process for the CPP

available and the need for extensive fine-tuning makes a predictor impossible to integrate into any scheduling solver.

#### 6.2.4. CONSTRUCTION AND FEASIBILITY OF A MONTHLY SOLUTION

As in Figure 3, we compare four approaches to construct monthly pairings. The first approach is a standard monthly solution, called “*Standard - initial*”, a “cyclic” weekly solution (from running the optimizer on a weekly CPP) rolled to cover the whole month (Desaulniers et al., 2020). A cyclic solution is where the number of crews in each city is the same at the beginning and end of the horizon. In the second approach, CNNs predict the flight-connection probabilities. Then, using the same heuristics as in Yaakoubi et al. (2020) (see Appendix D.3), we build a monthly crew pairing called “*CNN - initial*”. In the third and fourth approaches, CNN-CRF (Chu et al., 2016) and Struct-CKN are used to build monthly crew pairings called “*CNN-CRF - initial*”, and “*Struct-CKN - initial*”. Then, we break all illegal pairings and freeze the legal sub-part in initial crew pairings, resulting in “*Standard - feasible*”, “*CNN - feasible*”, “*CNN-CRF - feasible*” and “*Struct-CKN - feasible*”. We generate deadheads on all flights and pass it to the solver.

Note that to reproduce the following results, we obtain the commercial dataset from Yaakoubi et al. (2019), which cannot be distributed because it contains too much flight-data information sensitive to airlines’ operations. Since the first step of the CPP solver can solve up to several thousand flights, we are constrained to use two-day windows. Since pairings extend to over two days, it cannot “fix” the mispredicted and illegal pairings. Therefore, the more illegal pairings the constructed solution contains, the longer it will take the solver to find a suitable final solution. Furthermore,

Table 3. Characteristics of Monthly Solutions

	#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	<b>11.07</b>
Struct-CKN - feasible	<b>4 015</b>	16.46	

when the constructed solution is highly infeasible, it can only be proposed as initial clusters as in all past research (e.g., Yaakoubi et al. (2019; 2020)) and not as an initial solution. In this case, a generic standard initial solution is used, and since the initial solution and the initial clusters are different, an adaptation strategy is required to adapt the proposed clusters of the current window to the solution of the previous window. This is a major limitation of past research since the explored neighborhood needed to be large enough to reach a good LP (linear programming) solution but small enough to maintain a small number of fractional variables permitting to have an efficient heuristic branch-&-bound. Thus, not only can we conclude that the primary metric of interest in our case is the feasibility of the constructed monthly solution. But, it also becomes crucial to propose a feasible pairing solution that can be proposed both as initial clusters and as an initial solution, therefore bypassing the adaptation strategy and permitting to reduce the resolution time (by reducing the neighborhood in which to explore in order to find a suitable solution).

Table 3 summarizes the computational results on the feasibility and characteristics of constructed monthly pairings. First, breaking all illegal pairings in “*Standard - initial solution*” removes 50.56% of the pairings, while that in “*CNN - initial solution*” led to removing 21.05%. Note that even though the test error is low for CNNs, due to the large number of infeasible pairings, running the optimization using this initial solution is problematic. Breaking all illegal pairings in “*CNN-CRF - initial solution*” removes 12.12% of the pairings, while only 11.07% are removed from “*Struct-CKN - initial solution*”. Clearly, although Struct-CKN has 97% fewer parameters than other methods and its hyper-parameters are not exhaustively fine-tuned, it outperforms other methods in terms of test error and feasibility. In what follows, we provide the constructed monthly pairings to the solver and compare the resulting solutions when using the baseline solution (Desaulniers et al., 2020), CNN (Yaakoubi et al., 2020) and the proposed approach (Struct-CKN).



Table 4. Computational Results for Monthly Solutions

	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	<b>18.68</b>	<b>4.20</b>	<b>915</b>	64.48
CNN+	18.62	3.34	997	126.62
Struct-CKN+	<b>17.15</b>	<b>0.59</b>	<b>583</b>	41.44

### 6.2.5. RESULTS ON THE CREW PAIRING PROBLEM

As in Figure 3, the baseline solution (“Baseline”) for the monthly CPP is obtained by feeding the solver initial clusters from “Standard - feasible” (Desaulniers et al., 2020). The previous state-of-the-art was obtained by feeding the solver initial clusters from “CNN - feasible”, yielding a solution called “CNN”. Instead of CNNs, we can use Struct-CKN to propose **both initial clusters and an initial solution** from “Struct-CKN - feasible” to the solver, yielding a monthly solution called “Struct-CKN”. To work on finding the best monthly solution possible and overcome the limitations of using the windowing approach, we feed the solution obtained from “Struct-CKN” to the solver (again) as initial clusters and initial solution, yielding “Struct-CKN+”. Because we cannot use the constructed monthly pairing as an initial solution, when using CNNs, we claim that re-running the optimization using the solution “CNN” as initial clusters does not improve the monthly solution much. To support our claim, we feed solution “CNN” to the solver as initial clusters, yielding “CNN+” (see Figure 3). See Appendix D.5 for detailed statistics of the optimization process.

Table 4 reports computational results for the final monthly solution. Note that we do not report variances since the CPP solver is deterministic. Struct-CKN outperforms both Baseline and CNN reducing the solution cost (in millions of dollars) and cost of global constraints by 9.51% and 80.25%, respectively, while also being 33% faster than CNN. By re-running the optimization, on the one hand, CNN+ does not improve the solution much. On the other hand, the Struct-CKN+ solution yields the best statistics, reducing solution cost and cost of global constraints by 16.93% and 97.24%, respectively. More interestingly, the number of deadheads (see Section 6.2.1) is reduced by 41.23%, compared to Baseline. Therefore, we can conclude that proposing a feasible monthly solution both as initial clusters and as an initial solution allowed us to achieve better results in less time and to provide the possibility to re-optimize the solution and improve it further. Thus, this permits to update the training solutions, suggesting that Struct-CKN can be further optimized and that further research to avoid the windowing approach and use a one-month window can present better results than the current version of solver.

## 7. Conclusion

Seeking an initial solution to a crew pairing solver, this study proposes Struct-CKN, a new deep structured predictor. Its supervised use outperforms state-of-the-art methods in terms of primal sub-optimality of the structured prediction “layer” and test accuracy on the OCR dataset. The proposed method is then applied on a flight-connection dataset, modeled as a general CRF capable of incorporating local constraints in the learning process. To warm-start the solver, we use Struct-CKN to propose initial clusters and an initial solution to the solver, reducing the solution cost by 17% (a gain of millions of dollars) and the cost of global constraints by 97%, compared to baselines. Future research will look into combining deep structured methods with various operations research methods and designing new reactive/learning metaheuristics that learn to guide the search for better solutions in real-time.

## Acknowledgements

We are thankful to the anonymous reviewers and the meta-reviewer for their valuable comments that improved the quality of this work. We also would like to thank Andjela Mladenovic, Gina Arena, Joey Bose, Mehdi Abbana Bennani, and Stephanie Cairns for their constructive comments regarding this work. We thank Iban Harlouchet for his involvement during the first phase of the project. This work was supported by IVADO, a Collaborative Research and Development Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC), by the Canada CIFAR AI Chair Program, and AD OPT, a division of IBS Software. We would like to thank these organizations for their support and confidence. Simon Lacoste-Julien is a CIFAR Associate Fellow of the Learning in Machines & Brains program.

## References

- Belanger, D. and McCallum, A. Structured prediction energy networks. In *International Conference on Machine Learning*, 2016.
- Bengio, Y., Lodi, A., and Prouvost, A. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*, 2020.
- Bietti, A. and Mairal, J. Group invariance, stability to deformations, and complexity of deep convolutional representations. *JMLR*, 2019.
- Bottou, L. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*. Springer, 2012.

- Buchta, C., Kober, M., Feinerer, I., and Hornik, K. Spherical k-means clustering. *Journal of Statistical Software*, 2012.
- Chandra, S., Usunier, N., and Kokkinos, I. Dense and low-rank gaussian crfs using deep embeddings. In *IEEE International Conference on Computer Vision*, 2017.
- Chen, L.-C., Schwing, A., Yuille, A., and Urtasun, R. Learning deep structured models. In *International Conference on Machine Learning*, 2015.
- Chu, X., Ouyang, W., Li, h., and Wang, X. Crf-cnn: Modeling structured information in human pose estimation. In *Advances in Neural Information Processing Systems 29*, pp. 316–324. Curran Associates, Inc., 2016.
- Desaulniers, G., Desrosiers, J., Dumas, Y., Marc, S., Rioux, B., Solomon, M. M., and Soumis, F. Crew pairing at Air France. *European journal of operational research*, 97(2), 1997.
- Desaulniers, G., Lessard, F., Mohammed, S., and François, S. Dynamic constraint aggregation for solving very large-scale airline crew pairing problems. *Les Cahiers du GERAD*, G–2020(21), 2020.
- Elhallaoui, I., Metrane, A., Soumis, F., and Desaulniers, G. Multi-phase dynamic constraint aggregation for set partitioning type problems. *Mathematical Programming*, 123(2), 2010.
- Gasse, M., Chételat, D., Ferroni, N., Charlin, L., and Lodi, A. Exact combinatorial optimization with graph convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2019.
- Graber, C., Meshi, O., and Schwing, A. Deep structured prediction with nonlinear output transformations. In *Advances in Neural Information Processing Systems*, 2018.
- Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., and Schmidhuber, J. Lstm: A search space odyssey. *IEEE transactions on neural networks and learning systems*, 28(10), 2016.
- Gygli, M., Norouzi, M., and Angelova, A. Deep value networks learn to evaluate and iteratively refine structured outputs. *International Conference on Machine Learning*, 2017.
- Kasirzadeh, A., Saddoune, M., and Soumis, F. Airline crew scheduling: models, algorithms, and data sets. *EURO Journal on Transportation and Logistics*, 2017.
- Khalil, E., Dai, H., Zhang, Y., Dilkina, B., and Song, L. Learning combinatorial optimization algorithms over graphs. In *Advances in Neural Information Processing Systems*, 2017.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2016.
- Kumar, S., August, J., and Hebert, M. Exploiting inference for approximate parameter learning in discriminative fields: An empirical study. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*. Springer, 2005.
- Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. Block-coordinate frank-wolfe optimization for structural svms. *International Conference on Machine Learning*, 2012.
- Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. Block-coordinate frank-wolfe optimization for structural svms. *International Conference on Machine Learning*, 2013.
- Le Priol, R., Piché, A., and Lacoste-Julien, S. Adaptive stochastic dual coordinate ascent for conditional random fields. *Conference on Uncertainty in Artificial Intelligence*, 2018.
- Leblond, R., Alayrac, J.-B., Osokin, A., and Lacoste-Julien, S. Searnn: Training rnns with global-local losses. *International Conference on Machine Learning*, 2017.
- Mairal, J. End-to-end kernel learning with supervised convolutional kernel networks. In *Advances in Neural Information Processing Systems*, 2016.
- Mairal, J., Koniusz, P., Harchaoui, Z., and Schmid, C. Convolutional kernel networks. *Advances in Neural Information Processing Systems*, 2014.
- Martins, A. F., Figueiredo, M. A., Aguiar, P. M., Smith, N. A., and Xing, E. P. Ad3: Alternating directions dual decomposition for map inference in graphical models. *Journal of Machine Learning Research*, 2015.
- Müller, A. C. and Behnke, S. Pystruct: learning structured prediction in python. *J. Mach. Learn. Res.*, 15, 2014.
- Murphy, K. P. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Owerko, D., Gama, F., and Ribeiro, A. Optimal power flow using graph neural networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style,

- high-performance deep learning library. *Advances in Neural Information Processing Systems* 32, 2019.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2011.
- Pérez-Cruz, F., Ghahramani, Z., and Pontil, M. Kernel conditional graphical models. In *Predicting structured data*, 2007.
- Scharstein, D. and Pal, C. Learning conditional random fields for stereo. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007.
- Schmidt, M., Babanezhad, R., Ahmed, M., Defazio, A., Clifton, A., and Sarkar, A. Non-uniform stochastic average gradient method for training conditional random fields. In *AISTATS*, 2015.
- Shalev-Shwartz, S. and Zhang, T. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 2013.
- Shalev-Shwartz, S. and Zhang, T. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *International Conference on Machine Learning*, 2014.
- Taskar, B., Guestrin, C., and Koller, D. Max-margin markov networks. In *Advances in Neural Information Processing Systems*, 2004.
- Tran, S. N., Garcez, A. d., Weyde, T., Yin, J., Zhang, Q., and Karunanithi, M. Sequence classification restricted boltzmann machines with gated units. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- Yaakoubi, Y. *Combiner intelligence artificielle et programmation mathématique pour la planification des horaires des équipages en transport aérien*. PhD thesis, Polytechnique Montréal, 2019.
- Yaakoubi, Y., Soumis, F., and Lacoste-Julien, S. Flight-connection prediction for airline crew scheduling to construct initial clusters for OR optimizer. *Les Cahiers du GERAD*, G-2019(26), 2019.
- Yaakoubi, Y., Soumis, F., and Lacoste-Julien, S. 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.
- Yen, J. W. and Birge, J. R. A stochastic programming approach to the airline crew scheduling problem. *Transportation Science*, 40(1), 2006.
- Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., and Torr, P. H. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, 2015a.
- Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., Huang, C., and Torr, P. H. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, 2015b.
- Zhong, P. and Wang, R. Using combination of statistical models and multilevel structural information for detecting urban areas from a single gray-level image. *IEEE transactions on geoscience and remote sensing*, 45(5), 2007.

The appendix is organized as follows: we first provide details about the convolutional kernel networks (CKNs) and the multilayer convolutional kernel (MCK) in Appendix A. Next, we provide details on the Structured SVM approach in Appendix B.2. Then, we present additional results on the OCR dataset in Appendix C. Then, we present the prediction problem formulation in Appendix D.1 and describe the flight-connection dataset in Appendix D.2. Finally, we provide details about the optimization process and instances used in our paper in Appendix D.4 and present additional results for the Crew Pairing Problem (CPP) optimization in Appendix D.5. Note that in the appendix, we will refer to Struct-CKN using CRF models as Struct-CKN-SDCA. Instead of CRFs, we can also use SSVM predictors, in particular the block-coordinate Frank-Wolfe (BCFW) algorithm, which we will refer to as Struct-CKN-BCFW.

## A. Convolutional Kernel Networks

This section contains additional details on CKNs (Mairal et al., 2014). We first present the Kernel trick, the reproducing kernel Hilbert space, and the Nyström low-rank approximation in Sections A.1, A.2, and A.3, respectively. Then, Section A.4 presents MCKs.

### A.1. Kernel Trick

The kernel trick consists of embedding a “raw input space”  $\mathcal{X}$  into a high dimensional Hilbert “feature space”  $\mathcal{H}$ , through a possibly non-linear mapping  $\varphi : \mathcal{X} \rightarrow \mathcal{H}$ , where the inner product in  $\mathcal{H}$  admits the computation formula as in (8), with  $K : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ .

$$\langle \varphi(x), \varphi(x') \rangle_{\mathcal{H}} = K(x, x'), \quad x, x' \in \mathcal{X} \quad (8)$$

The “kernel function”  $K$  is thought of as specifying similarity between elements of  $\mathcal{X}$ . In contrast, the similarities in  $\mathcal{H}$  are expressed as simple inner products. The choice of an appropriate  $\varphi$  can result in making the image of the data subset of  $\mathcal{X}$  quasi-linearly separable, in the sense that the image of  $\mathcal{X}$  and the image of another subset disjoint to  $\mathcal{X}$  could (almost) be separated with a linear predictor. When a linear predictor on  $\mathcal{H}$  only uses inner products of  $\mathcal{H}$  elements, then the formula (8) makes the training of this linear predictor computationally feasible.

The next section reviews a specific class of feature spaces that are built from positive definite kernels.

### A.2. Reproducing Kernel Hilbert Space

**Definition 1 (Positive definite kernels)** A positive definite (p.d.) kernel on a set  $\mathcal{X}$  is a function  $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  that is symmetric (i.e.  $K(x, x') = K(x', x), \forall x, x' \in \mathcal{X}$ ), and for which, for any  $N \in \mathbb{N}$  and any  $(x_1, \dots, x_N) \in \mathcal{X}^N$ , the Gram matrix (or similarity matrix)  $[K]_{ij} = K(x_i, x_j)$  is positive semi-definite (which means that for any  $(a_1, \dots, a_N) \in \mathbb{R}^n$ ,  $\sum_{i=1}^N \sum_{j=1}^N a_i a_j K(x_i, x_j) \geq 0$ ).

**Theorem 1 (RKHS (Bietti & Mairal, 2019))** Let  $K$  be a p.d. kernel on a set  $\mathcal{X}$ . There exists a unique Hilbert space  $\mathcal{H} \subset \mathcal{X}^{\mathbb{R}}$  such that

1.  $\{K : y \mapsto K(x, y), \mathcal{X} \mapsto \mathcal{R}\}_{x \in \mathcal{X}} \subset \mathcal{H}$
2.  $\forall f \in \mathcal{H}, \forall x \in \mathcal{X}, f(x) = \langle f, K \rangle_{\mathcal{H}}$

In particular, the application

$$\begin{aligned} \varphi : \mathcal{X} &\mapsto \mathcal{H} \\ x &\mapsto K \end{aligned}$$

which maps  $\mathcal{X}$  in the “feature space”  $\mathcal{H}$  satisfies

$$\langle \varphi(x), \varphi(x') \rangle_{\mathcal{H}} = K(x, x') \quad x, x' \in \mathcal{X} \quad (9)$$

We give now two classical examples of RHKSs associated to the euclidean space  $\mathcal{X} = \mathbb{R}^d$ : the linear kernel and the Gaussian kernel. The linear kernel on  $\mathcal{X}$  is defined to be :

$$K_{\text{lin}}(x, x') = \langle x, x' \rangle_{\mathbb{R}^d} \quad (10)$$

The RKHS associated to  $(\mathcal{X}, K_{lin})$  is the Hilbert space  $\mathcal{H} = \{f_w = \langle \bullet, w \rangle_{\mathbb{R}^d}\}_{w \in \mathbb{R}^d}$  endowed with the inner product  $\langle f_w, f_v \rangle_{\mathcal{H}} = \langle w, v \rangle_{\mathbb{R}^d}$ . This RKHS is isometric to  $\mathbb{R}^d$ . The Gaussian Kernel with bandwidth  $\sigma$  on  $\mathcal{X}$  is defined to be

$$K_{\text{Gauss}}(x, y) = e^{-\frac{\|x-y\|^2}{2\sigma^2}} \quad (11)$$

The RKHS  $\mathcal{H}$  associated to  $(\mathcal{X}, K_{\text{Gauss}})$  is an infinite-dimensional feature space, and all points of  $\mathcal{X}$  are mapped to the unit sphere of  $\mathcal{H}$  ( $\|\varphi(x)\|_{\mathcal{H}}^2 = K(x, x) = 1$ ).

**Theorem 2 (Representer theorem)** *Let  $\mathcal{X}$  be a set endowed with a p.d. kernel  $K$ , let  $\mathcal{H}$  be its RKHS and  $S = \{x_1, \dots, x_n\} \subset \mathcal{X}$  a training data set. If  $\Psi : \mathbb{R}^{n+1} \mapsto \mathbb{R}$  is strictly increasing with respect to the last variable, then*

$$\begin{aligned} f^* &= \arg \min_{f \in \mathcal{H}} \Psi(f(x_1), \dots, f(x_n), \|f\|_{\mathcal{H}}) \\ \implies f^* &\in \text{Span}(K_{x_1}, \dots, K_{x_n}) \end{aligned}$$

Remark:  $f^*$  lives in a subspace of dimension  $n$ , although  $\mathcal{H}$  can be of infinite dimension.

### A.3. Nyström Low-rank Approximation

Consider a p.d. kernel  $K : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$  and its RKHS  $\mathcal{H}$ , with the mapping  $\varphi : \mathcal{X} \mapsto \mathcal{H}$  such that  $K(x, x') = \langle \varphi(x), \varphi(x') \rangle_{\mathcal{H}}$ . The Nyström approximation consists in replacing any point  $\varphi(x)$  in  $\mathcal{H}$  by its orthogonal projection  $\Pi_{\mathcal{F}}(x)$  onto a finite-dimensional subspace

$$\mathcal{F} := \text{Span}(f_1, \dots, f_p) \text{ with } p \ll n \quad (12)$$

where the  $f'_i$  are *anchor points* in  $\mathcal{H}$  (defined below).

This projection is equivalent to

$$\Pi_{\mathcal{F}}(x) := \sum_{j=1}^p \beta_j^*(x) f_j \quad (13)$$

with

$$\beta^*(x) = \arg \min_{\beta \in \mathbb{R}^p} \left\| \varphi(x) - \sum_{j=1}^p \beta_j f_j \right\|_{\mathcal{H}}^2 \quad (14)$$

Noting  $[K_f]_{jl} = \langle f_j, f_l \rangle_{\mathcal{H}}$  and  $f(x) = (f_1(x), \dots, f_p(x)) \in \mathbb{R}^p$ , it is quickly checked that  $\beta^*(x) = K_f^{-1} f(x)$ , where  $K_f^{-1}$  is the inverse of the kernel matrix  $K_f$  (or its pseudo-inverse, when the matrix  $K_f$  is not full rank). Thus,

$$\varphi(x) \approx \sum_{j=1}^p \beta_j^*(x) f_j \quad (15)$$

and

$$\langle \varphi(x), \varphi(x') \rangle_{\mathcal{H}} \approx \beta^*(x)^\top K_f \beta^*(x') \quad (16)$$

Defining the mapping

$$\begin{aligned} \Gamma : \mathcal{X} &\mapsto \mathbb{R}^p \\ x &\mapsto \Gamma(x) = K_f^{-1/2} f(x) \end{aligned}$$

as  $\Gamma(x) = K_f^{-1/2} \beta^*(x)$ , we see that

$$K(x, x') \approx \langle \Gamma(x), \Gamma(x') \rangle_{\mathbb{R}^p}$$

Consequently, the mapping  $\varphi : \mathcal{X} \mapsto \mathcal{H}$  is approximated by a mapping  $\Gamma : \mathcal{X} \mapsto \mathbb{R}^p$  such that

$$\langle \varphi(x), \varphi(x') \rangle_{\mathcal{H}} \approx \langle \Gamma(x), \Gamma(x') \rangle_{\mathbb{R}^p}$$

There exists various methods for choosing the anchor points  $f_j$ 's. When  $\mathcal{X} = \mathbb{R}^d$ , one such method consists in performing a  $K$ -means algorithm on a training data set  $S = \{x_1, \dots, x_n\} \subset \mathcal{X}$ , to obtain  $p$  centroids  $\mathbf{z}_1, \dots, \mathbf{z}_{p_1} \in \mathbb{R}^d$ , and define anchor points as  $f_j = \varphi(\mathbf{z}_j)$ , for  $j = 1, \dots, p$ .

#### A.4. Multilayer Convolutional Kernel (Mairal, 2016)

Two principal advantages of a CNN over a fully-connected neural network are: (i) sparsity - each nonlinear convolutional filter acts only on a local patch of the input, (ii) parameter sharing - the same filter is applied to each patch.

**Definition 2 (Image feature map and patch feature map)** *An image feature map  $\varphi$  is a function  $\varphi : \Omega \rightarrow \mathcal{H}$ , where  $\Omega$  is a discrete subset of  $[0, 1]^2$  representing a set of pixel locations, and  $\mathcal{H}$  is a Hilbert space representing a feature space. A patch feature map is the same as an image feature map with  $\Omega$  replaced by a patch of pixels  $\mathcal{P}$  centered at 0.*

**Definition 3 (Convolutional kernel)** *Let us consider two images represented by two image feature maps, respectively  $\varphi$  and  $\varphi' : \Omega \rightarrow \mathcal{H}$ . The one-layer convolutional kernel between  $\varphi$  and  $\varphi'$  is defined as*

$$K(\varphi, \varphi') = \sum_{z \in \Omega} \sum_{z' \in \Omega} \left\{ \|\varphi(z)\|_{\mathcal{H}} \|\varphi'(z')\|_{\mathcal{H}} e^{-\frac{1}{2\beta^2} \|z-z'\|_2^2} e^{-\frac{1}{2\sigma^2} \|\tilde{\varphi}(z) - \tilde{\varphi}'(z')\|_2^2} \right\}$$

where  $\tilde{\varphi}(z) = (1/\|\varphi(z)\|_{\mathcal{H}}) \varphi(z)$  if  $\varphi(z) \neq 0$  and  $\tilde{\varphi}(z) = 0$  otherwise. Idem for  $\tilde{\varphi}'(z)$ .

Remark 1: The role of  $\beta$  is to control how much the kernel is locally shift-invariant.

Remark 2:  $K$  is built on a kernel  $k$  on  $\mathcal{H}$ , of the form

$$k(h, h') = \|h\|_{\mathcal{H}} \|h'\|_{\mathcal{H}} \kappa \left( \frac{\langle h, h' \rangle}{\|h\|_{\mathcal{H}} \|h'\|_{\mathcal{H}}} \right), \quad h, h' \in \mathcal{H} \quad (17)$$

According to a classical result (Schoenberg, 1942), when  $\kappa$  is smooth with non-negative Taylor expansion coefficients,  $k$  is a p.d. kernel. Consequently,  $K$  is also a p.d. kernel.

**Definition 4 (Multilayer convolutional kernel)** *Let us consider a discrete set of pixel locations  $\Omega_{k-1} \subseteq [0, 1]^2$  and a Hilbert space  $\mathcal{H}_{k-1}$ . We build a new set  $\Omega_k$  and a new Hilbert space  $\mathcal{H}_k$  as follows:*

(i) *choose a patch shape  $\mathcal{P}_k$  defined as a discrete symmetric subset of  $[-1, 1]^2$ , and a set of pixel locations  $\Omega_k$  such that for all location  $\mathbf{z}_k$  in  $\Omega_k$ , the patch  $\{\mathbf{z}_k\} + \mathcal{P}_k$  is a subset of  $\Omega_{k-1}$ ; In other words, each pixel location  $\mathbf{z}_k$  in  $\Omega_k$  corresponds to a valid patch of pixel locations in  $\Omega_{k-1}$  centered at  $\mathbf{z}_k$ .*

(ii) *define the convolutional kernel  $K_k$  on the “patch” feature maps  $\mathcal{P}_k \rightarrow \mathcal{H}_{k-1}$ , by replacing in (17):  $\Omega$  by  $\mathcal{P}_k$ ,  $\mathcal{H}$  by  $\mathcal{H}_{k-1}$ , and  $\sigma, \beta$  by appropriate smoothing parameters  $\sigma_k, \beta_k$ . We denote by  $\mathcal{H}_k$  the Hilbert space for which the positive definite kernel  $K_k$  is reproducing.*

*An image represented by a feature map  $\varphi_{k-1} : \Omega_{k-1} \rightarrow \mathcal{H}_{k-1}$  at layer  $k-1$  is now encoded in the  $k$ -th layer as  $\varphi_k : \Omega_k \rightarrow \mathcal{H}_k$ , where for all  $\mathbf{z}_k$  in  $\Omega_k$ ,  $\varphi_k(\mathbf{z}_k)$  is the representation in  $\mathcal{H}_k$  of the patch feature map  $\mathbf{z} \mapsto \varphi_{k-1}(\mathbf{z}_k + \mathbf{z})$  for  $\mathbf{z}$  in  $\mathcal{P}_k$ .*

In other words, given  $\varphi_{k-1} : \Omega_{k-1} \rightarrow \mathcal{H}_{k-1}$ , we first define  $\Omega_k$  and  $\mathcal{P}_k$  such that for any  $\mathbf{z}_k \in \Omega_k$ ,  $\mathbf{z}_k + \mathcal{P}_k \subset \Omega_{k-1}$ . Second, we define  $\mathcal{H}_k$  as being the RKHS of the set of patch feature maps

$$\{\mathbf{z} \mapsto \varphi(\mathbf{z}_k + \mathbf{z}), \mathcal{P}_k \mapsto \mathcal{H}_{k-1} : \varphi : \Omega_{k-1} \rightarrow \mathcal{H}_{k-1}, \mathbf{z}_k \in \Omega_k\}$$

endowed with the positive definite kernel  $K_k$  adapted from (17), and we set  $\varphi_k(\mathbf{z}_k)$  as being the representation in  $\mathcal{H}_k$  of the patch feature map  $\mathbf{z} \mapsto \varphi_{k-1}(\mathbf{z}_k + \mathbf{z})$ . Thus we obtain  $\varphi_k : \Omega_k \rightarrow \mathcal{H}_k$ , and for any  $\varphi_{k-1}, \varphi'_{k-1} : \Omega_{k-1} \rightarrow \mathcal{H}_{k-1}$ , we have  $K_k(\varphi_{k-1}, \varphi'_{k-1}) = \langle \varphi_k, \varphi'_k \rangle_{\mathcal{H}_k}$ .

## B. Graph-Based Learning

### B.1. Conditional Random Fields

Algorithm 2 describes this variant of SDCA, biasing the sampling towards examples with large duality gaps  $g$ . If we assume that the graph has a junction tree structure  $\mathbb{T} = (\mathbb{C}, \mathbb{S})$ , where  $\mathbb{C}$  is the set of maximal cliques and  $\mathbb{S}$  the set of separators, we can run message passing on the junction tree to infer new marginals given weights  $w$ :  $\hat{\mu}_i(w) = p(y_C = \cdot | x_i; w)$ , and recover joint probability  $\alpha_i(y)$  as function of marginals  $\mu_{i,C}$ :  $\alpha_i(y) = \frac{\prod_{C \in \mathbb{C}} \mu_{i,C}(y_C)}{\prod_{S \in \mathbb{S}} \mu_{i,S}(y_S)}$ . This allows computing the entropy as

in (18) and the Kullback-Leibler divergence of joints as in (19), using only the marginals  $\mu_i$  and  $\nu_i$  of  $\alpha_i$  and  $\beta_i$ , respectively. With (18), we compute the dual objective value, and perform the line search, as in (20).

$$\tilde{H}(\mu_i) = H(\alpha_i) = \sum_C H(\mu_{i,C}) - \sum_S H(\mu_{i,S}) \quad (18)$$

$$\tilde{D}(\mu_i|\nu_i) = \sum_C D_{KL}(\mu_{i,C}|\nu_{i,C}) - \sum_S D_{KL}(\mu_{i,S}|\nu_{i,S}) \quad (19)$$

$$\gamma^* = \operatorname{argmax}_{\gamma \in [0,1]} \tilde{H}(\mu_i^{(t)} + \gamma \delta_i) - \frac{\lambda n}{2} \|w^{(t)} + \gamma v_i\|^2 \quad (20)$$

---

**Algorithm 2** SDCA for CRF (Le Priol et al., 2018)

- 1: Initialize  $\mu_i^{(0)} \in \Pi_C^{\Delta_C}, \forall i; w^{(0)} = \hat{w}(\mu^{(0)}) = \frac{1}{\lambda n} \sum_i B_i \mu_i^{(0)}$ , where  $B_i$  is the horizontal concatenation of the feature vectors over cliques
  - 2: (Optional) Initialize the duality gaps  $g_i = 100, \forall i$
  - 3: **for**  $t=1 : n_{Ep\_SDCA}$  **do**
  - 4:   Sample  $i$  uniformly in  $\{1, \dots, n\}$
  - 5:   (Alternatively) Sample  $i$  with  $P(i) \propto g_i$
  - 6:   Let  $\nu_{i,C}(y_C) = p(y_C|x_i; w^{(t)}), \forall C \in \mathbb{C}$
  - 7:   (Optional) Let  $g_i = \tilde{D}(\mu_i|\nu_i)$  (19)
  - 8:   Let ascent direction be :  $\delta_i = \nu_i - \mu_i^t$ , and the primal direction be :  $v_i = \frac{1}{\lambda n} \hat{w}(\delta_i)$
  - 9:   Use line search to get the optimal step size (20)
  - 10:   Update  $\mu_i^{(t+1)} = \mu_i^{(t)} + \gamma^* \delta_i$ , and  $w^{(t+1)} = \hat{w}(\mu^{(t+1)}) = w^{(t)} + \gamma^* v_i$
  - 11: **end for**
- 

Note that as in Section 5, to do inference for CRF models, we use AD3 (Martins et al., 2015) for approximate maximum a posteriori (MAP) inference. SDCA needs a marginalization oracle, whereas AD3 is used to do an approximate MAP. In order to use SDCA with AD3, we propose a simple approximation, using MAP label estimates. We first run an approximate MAP inference algorithm and then imagine that the marginals put a unit probability at the approximate MAP solution and zero elsewhere (Kumar et al., 2005; Zhong & Wang, 2007; Scharstein & Pal, 2007). This is a heuristic method, but it can be expected to work well when the estimated MAP solution is close to the true MAP, and the conditional distribution is strongly “peaked”.

## B.2. Structural SVM - Block-Coordinate Frank-Wolfe

To train CRFs, we either use a Structured SVM solver using BCFW algorithm (Lacoste-Julien et al., 2013), or a stochastic optimization algorithm using SDCA (Le Priol et al., 2018). In this Section, we briefly review SSVM.

Given a training set of input-output structure pairs  $\{(x_1, y_1), \dots, (x_n, y_n)\} \in \mathcal{X} \times \mathcal{Y}$ , a loss function  $\Delta : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$  and a joint feature map  $F : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  that encodes the input/output information necessary for prediction, we aim to learn  $w$  so that the prediction function  $f(x) = \operatorname{argmax}_y \langle w, F(x, y) \rangle$  minimizes the expected loss on future data. To do so, we use SSVM predictors derived by maximum margin framework, and enforce correct output to be better than others by a margin:  $\langle w, F(x_n, y_n) \rangle \geq \Delta(y_n, y) + \langle w, F(x_n, y) \rangle; \forall y \in \mathcal{Y}$ . This is a convex optimization problem, but non-differentiable, with many equivalent formulations, resulting in different training algorithms. The Frank-Wolfe algorithm, one of such training algorithms, considers the convex minimization problem  $\min_{\alpha \in M} f(\alpha)$ , where  $M$  is compact, and  $f$  is continuously differentiable. It only requires optimizing linear functions over  $M$ . Since we experiment on large datasets, we use block-coordinate Frank-Wolfe (BCFW) algorithm (Lacoste-Julien et al., 2012) showing better convergence than the original or the batch version.

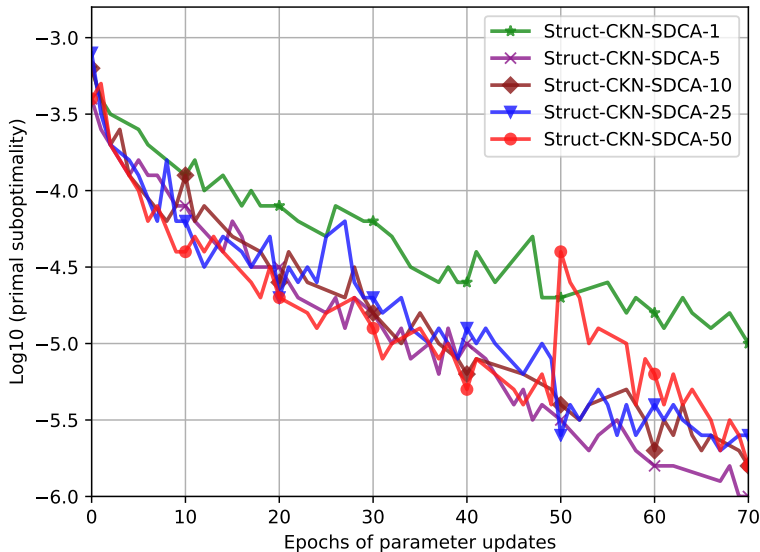


Figure 4. Struct-ckn With Gap Sampling With Various  $n_{Ep\_SDCA}$  Values, as Indicated by the Number in the Legend

Table 5. Test Error on the OCR Dataset

	SDCA - linear features (Le Priol et al., 2018)	LSTM (Greff et al., 2016)	CNN- CRF (Chu et al., 2016)	SCRBM (Tran et al., 2020)	Struct-CKN- BCFW	Struct-CKN- SDCA
Test error (%)	12.0	4.6	4.5	4.03	3.42	<b>3.40</b>

### C. Additional results for the Struct-CKN predictor on the OCR dataset

The hyperparameters for CKN are the number of layers, the number of filters and the size of patches. We use one layer and perform non-exhaustive grid search to set the number of filters (200) and the size of patches ( $5 \times 5$ ). We use cross-validation on the training set to measure the quality of the configuration and use different months for different samples to simulate the more realistic scenario in which we make a forecast over a new period. To run SDCA, we use  $\lambda = \frac{1}{n}$ , where  $n$  is the number of data points (sequences). As described in Algorithm 2, we use the gap sampling strategy: we sample proportionally to our current estimate of the duality gaps, with 80% of uniform sampling. Finally, we use the Newton-Raphson algorithm on the derivative of the line search function. We train the Struct-CKN-SDCA predictor using different values for  $n_{Ep\_SDCA}$ . As shown in Figure 4, we obtain similar results if we use values for  $n_{Ep\_SDCA}$  between 5 and 25. Performing a single update of the SDCA parameter before updating the CKN weights makes the learning process difficult and using a value of 50 leads to the optimizer overfitting, making the learning process more difficult afterwards.

**Oracle Calls.** Schmidt et al. (2015) compared the algorithms based on the number of oracle calls. This metric was suitable for the methods they compared. Both OEG and SAG-NUS use a line search where they call an oracle on each step. SDCA, SDCA-GAP and Struct-CKN-SDCA do not require the oracle to perform their line search. However, the oracle is message passing on a junction tree. It has a cost proportional to the size of the marginals. Each iteration of the line search requires computing the entropy of these marginals, or their derivatives. These costs are roughly the same.

**Parameter Updates.** To give a different perspective, we also report the log of the sub-optimality against the number of parameter updates. This removes the additional cost of the line search for all methods.

**Test Error on the OCR Dataset.** We report in Table 5 the test errors on the OCR dataset.



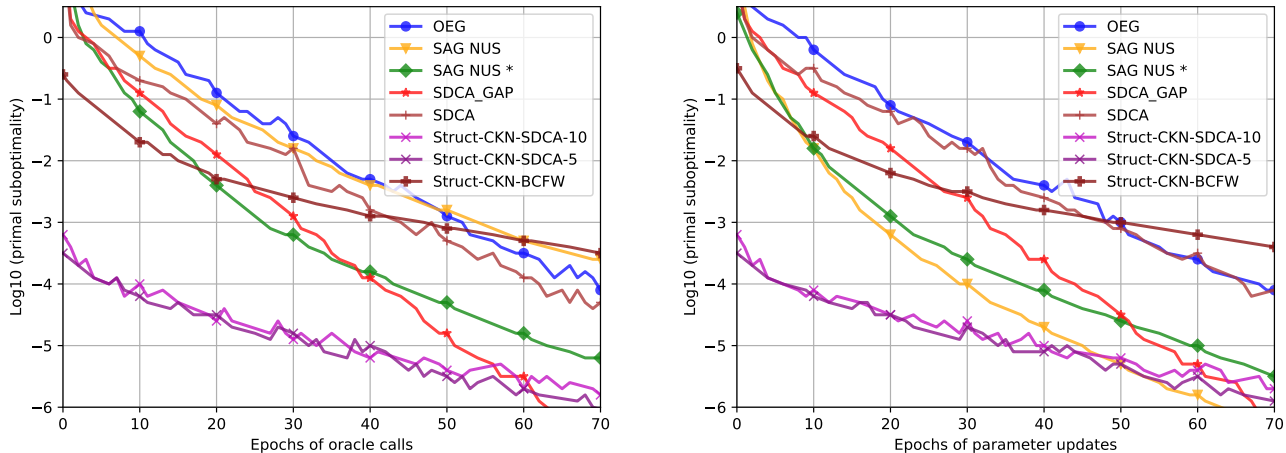


Figure 5. Comparison of Primal Sub-optimality According to the Number of Oracle Calls (Left) or Parameter Updates (Right). SDCA Refers to Uniform Sampling. SDCA-GAP Refers to Sampling Gap Sampling 80% of the Time. SAG-NUS Performs a Line Search at Every Iteration. SAG-NUS\* Implements a Line-search Skipping Strategy. Struct-CKN-BCFW and Struct-CKN-SDCA Use BCFW and SDCA Respectively as Structured Predictors. Struct-CKN-SDCA-5 and Struct-CKN-SDCA-10 Use 5 and 10 Epochs Respectively to Train the SDCA Weights for Each Iteration.

### D. Crew Pairing Problem

This section contains additional details on the Crew Pairing Problem (CPP). We first describe the prediction problem formulation and the flight-connection dataset in Sections D.1 and D.2, respectively. We describe the construction of the monthly crew pairing, once the ML predictor is trained, in Section D.3. Then, we present in Section D.4 in full the optimization process used to solve the CPP, once we obtain the initial solution and initial clusters using the ML predictor. Finally, we report additional computational results for the CPP optimization per window and for monthly solution.

#### D.1. Prediction Problem Formulation

The classification problem is the following: given the information about an incoming flight in a specific connecting city, choose among all the possible departing flights from this city (which can be restricted to the next 48 hours), the one that the crew should follow. These departing flights will be identified by their flight code (approximately 2,000 possible flight codes in the flight-connection dataset). Different flights may share the same flight codes in some airline companies, as flights performed multiple times a week usually use the same flight code. Nevertheless, a flight is uniquely identified by its flight code, departing city, and day, information that can be derived from the information on the incoming flight and the 48-hour window.

Each flight gives the following 5 features that we can use in our classification algorithm:

- city of origin and city of destination (~200 choices);
- aircraft type (5 types);
- duration of flight (in minutes)
- arrival time (for an incoming flight) or departure time (for a departing flight).

Given that the aircrew arrived at a specific airport at a given time, we can use *a priori* knowledge to define which flights are possible. For example, it is not possible to make a flight that starts ten minutes after the arrival, nor is it possible five days later. Furthermore, it is rare that the type of aircraft changes between flights since each aircrew is formed to use one or two types of aircraft at most. The reader is referred to Kasirzadeh et al. (2017) for further details on the likelihood of these scenarios. Note that even though we use the aircraft routing in the pre-processing steps, the classifier cannot know which next flight uses the same aircraft, as the prediction performance of the ML predictor is similar whether we use this information as a feature or not. The prediction problem is, therefore, not sensitive to this information and can do without it.

```

O T 1821 0 CC 1 CAT 1 () "a" "" _2_456_ 2018/08/09
2018/09/03 X 2018/08/10 2018/08/20 2018/08/29
{
RPT ;
ABC05914 ORD 1 02:10 BUF;
ABC06161 BUF 1 09:50 ORD;
RLS;
}
    
```

Figure 6. Data Format for Aircrew Pairings

To construct the flight-connection dataset, the following conditions are used, which must always be met for the next flight performed by the crew:

- The departure time of the next flight should follow the arrival time of the previous flight to the connecting city;
- The departure time of the next flight should not exceed 48 hours following the arrival time of the previous flight to the connecting city;
- The departure city of the next flight should be identical to the connecting city in the previous flight;
- The aircraft type should be the same. Indeed, crew scheduling is separable by crew category and aircraft type or family (Kasirzadeh et al., 2017).

For each incoming flight, all the departing flights in the next 48 hours are considered as a set. Then, only those that are feasible according to the masking constraints introduced in Yaakoubi et al. (2019); Yaakoubi (2019) are considered to filter this set. The maximum number of possible flights is limited to 20, as it is sufficient in the airline industry, then the task is to predict the rank of the true label among this set of departing flights.

## D.2. Flight-connection Dataset

Note that the instances in the flight-connection dataset (Yaakoubi et al., 2019) originate from an anonymous major airline and consist of six monthly crew pairing solutions for approximately 200 cities and approximately 50,000 flights per month. Each instance contains a one-month flight schedule, engine rotations, and crew pairings, as well as a list of airports and bases. This information is used to create input for the learning phase. The dataset consists of approximately 1,000 different source-destination pairs, 2,000 different flight codes, and pairings start from 7 different bases.

As shown in Figure 6, this entry describes a pairing that takes place every Tuesday, Thursday, Friday, and Saturday between August 9, 2018 and September 3, 2018, except for the dates: 08/10, 08/20, and 08/29. The pairing contains flights between Buffalo Niagara International Airport (BUF) and O’Hare International Airport (ORD). Apart from this, we know the distance between the two airports is 761 Km, from which we can extrapolate the duration of the flight, based on an average aircraft flight speed, which is equivalent to roughly 1 hour and 26 minutes.

Note that the flight-connection dataset (Yaakoubi et al., 2019) contains an embedded representation of the flights’ features and not the original features. An advantage of this approach is to anonymize further the dataset, an important condition to release the dataset in public domain. The dataset contains the output of the embedding layer, therefore an embedded representation of the input. For each incoming flight, the embedding layer (dimension  $n_d$ ) is used to construct a feature representation for each of the 20 possible next flights. The concatenated matrix (2D matrix) is a  $n_d \times 20$  input, enabling the use of convolutional architecture across time. These representations were obtained during the last training epoch with a small tuned learning rate. This way, the representation of a city is not unique throughout the whole dataset. However, two representations of the same city are still similar.

Furthermore, note that to use Struct-CKN on the flight-connection dataset, we had to contact the authors of Yaakoubi et al. (2019) in order to obtain the arcs (the connections between flights). This information is not part of the flight-connection dataset accessible in the public domain since disclosing it would easily permit the user to reconstruct the aeronautical map and identify the anonymous airline. The reader interested in reproducing our results on the flight-connection dataset is encouraged to contact authors of the paper or us. Finally, note that in contrast to the paper where the authors tackle a weekly

crew pairing problem (approximately 10,000), and thus where the test set consists of the flights within a period of one week, our manuscript tackles a monthly crew pairing (50,000 flights/month). Therefore, although our test set is part of the flight-connection dataset, it is different from the test set considered in [Yaakoubi et al. \(2019\)](#); [Yaakoubi \(2019\)](#).

### D.3. From CNN Flight-connection Probabilities to Monthly Crew Pairings ([Yaakoubi et al., 2020](#))

In what follows, we describe the construction of the monthly crew pairing once the CNN predictor is trained. As stated in Section 6.2.2, we use the same methods and heuristics as in [Yaakoubi et al. \(2020\)](#) in order to obtain “*CNN - initial*”.

Upon the finalization of the flights-connection prediction CNN model training, we can use the same architecture to solve two other prediction problems on the test set (50,000 flights): (i) predict if each of the scheduled flights is the beginning of a pairing or not; and (ii) predict whether each flight is performed after a layover or not. In reality, the three predictors share the same representation. To solve these independent classification problems, we sum the three prediction problems’ cross-entropy losses when learning, therefore performing a multi-output classification.

**Start in a Base.** While training the CNN predictor to recognize whether a flight is the beginning of a pairing or not, it is possible that it misclassifies that a flight departing from a non-base city is indeed the beginning of a pairing. It is imperative to correct such false predictions in order to avoid pairings starting away from the base. Even though it is possible to construct a predictor to which only flights departing from the bases are given, it is more efficient and robust to use all flights in the training step. That way, the predictor learns a better representation of the input.

**No Layover Below a Threshold.** While training the CNN predictor to recognize whether there is a layover or not between two flights, and since this decision is independent of finding the next flight, we can use a threshold on the number of hours between the previous and the next flight, below which it does not make sense to make a layover. This threshold should be defined considering previous solutions or by a practitioner.

To build the crew pairing, we use a greedy heuristic to build a crew pairing. Specifically, we consider each flight that the model predicts at the beginning of a pairing as a first flight. Given this incoming flight, we predict whether the crew is making a layover or not. In both cases, we consider the incoming flight and predict the next one. The pairing ends when the maximal number of days permitted per pairing is approached. We can use the above heuristic to construct a solution for the testing data, obtaining a monthly crew pairing that can be fed as initial clusters to the CPP solver. Unfortunately, if one flight in the pairing is poorly predicted, as the flights are predefined, the crew can finish its pairing away from the base. To correct the pairings, we define a heuristic in which all pairings where the crews finish away from the base are deleted.

Because the monthly problem is solved using a windowing approach with one-week windows and two days of overlap period, constructing an initial partition for the entire month and using the subset in each window to feed the solver can be a major flaw. Such initial partition will have many inconsistencies with the solution of the previous window, particularly during the overlap period, such as flights belonging to two different clusters. We propose to adapt the proposed clusters to the solution of the previous window using the heuristics described above to construct the clusters of the current window in accordance with the solution found for the previous window and any inconsistency with the previous window is avoided, so the proposed partition is adapted to the current resolution. In addition, instead of only considering the flights that the model predicts at the beginning of a pairing as a first flight, incomplete clusters from the solution of the previous window starting during the overlap period are completed.

### D.4. Optimization Process

As shown in Figure 7, upon the finalization of training the Struct-CKN predictor weights, we build a monthly solution that we can provide to the optimization process as an initial solution and as initial clusters. Indeed, the optimization process requires a standard initial solution that is usually a “cyclic” weekly solution rolled to cover the whole month. In this monthly solution, pairings containing flights that have disappeared or have changed hours are broken; new pairings are then built to cover these flights as well, along with new flights that have appeared.

Instead of using a standard initial solution called “*Standard - initial*”, we use the Struct-CKN predictor to construct the monthly solution. One can use a greedy heuristic to build the crew pairing. Specifically, we consider each flight that the model predicts at the beginning of a pairing as a first flight. Given this incoming flight, we predict the next one. The pairing ends when the maximal number of days permitted per pairing is approached. We can use the above heuristic to construct a solution for test data, obtaining a monthly solution that can be fed both as an initial solution and as an initial partition to the

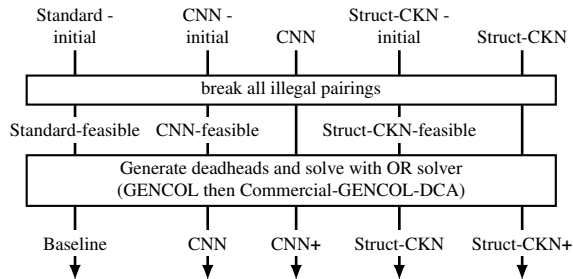


Figure 7. The Optimization Process for the CPP

solver. Unfortunately, if one flight in the pairing is poorly predicted, the crew can finish its pairing away from the base. Therefore, for all pairings where crews finish away from the base, we delete all flights performed after the last time that the crew arrives at the base. Note that pairings may begin before the bidding period to cover all active flights. Therefore, using the Struct-CKN predictor, we propose pairings on an extended period to cover the maximal number of flights. This leads to a monthly crew pairing schedule that we consider as an initial solution called “*Struct-CKN - initial - extended period*”. We can also restrict the pairings to be in the bid period. This leads to another monthly crew pairing that can also be considered as an initial solution called “*Struct-CKN - initial - bid period*”. Note that we only consider “*Struct-CKN - initial*” in the main paper. Then, we freeze the legal sub-part of pairings contained in the initial crew pairings by breaking all illegal pairings. Note that pairings starting before the bidding period are considered illegal since we would like to prevent the solver from generating pairings before the bid period starts. By breaking all illegal pairings from “*Standard - initial*”, we obtain a monthly solution called “*Standard - feasible*”. Similarly, by breaking all illegal pairings either from “*Struct-CKN - initial - extended period*” or “*Struct-CKN - initial - bid period*”, we obtain the same monthly solution called “*Struct-CKN - feasible*”.

To cover the flights that just became uncovered, we generate deadheads on all flights and use the GENCOL solver to optimize with legal sub-blocks imposed covering open (active) flights. In this approach, the problem is solved by a “rolling time horizon” approach. Because the GENCOL solver is able to solve up to a few thousand flights per window, we are constrained to use two-days windows. This means that the month is divided into overlapping time slices of equal length. Then, a solution is constructed greedily in chronological order by solving the problem *restricted* to each time slice sequentially, taking into account the solution of the previous slice through additional constraints. The monthly solution obtained is then passed through to the CPP solver. The clusters can either be extracted from this solution (as is usually done) or given separately (as is the case in our approach). Therefore, by providing an initial solution, we can not only expedite the optimization process, calculate the number of pairings considered legal, the number of flights covered but also propose clusters that are similar to the initial solution, thus reducing the degree of incompatibility between the current solution and the proposed pairings. Using the standard approach, we propose “*Standard - feasible*” as an initial solution to GENCOL and obtain a solution called “*GENCOL init.*”. Similarly, using the Struct-CKN predictor, we propose “*Struct-CKN - feasible*” as an initial solution to GENCOL and obtain a solution called “*Struct-CKN + GENCOL*”.

## D.5. Additional Results for The Crew Pairing Problem Optimization

In what follows, we include additional results on the flight-connections dataset, and the characteristics of the monthly solutions constructed using probabilities yielded by different predictors in Table 6 and Table 7, respectively. By using the monthly crew pairings to start the CPP solver, we report in Table 8 and 9 the computational results per window for the optimization process, and computation results for the final monthly solution, respectively.

We report in Table 6 the test error on the flight-connection dataset using (1) CNNs and Gaussian Process (GP) to search for the best configuration of hyperparameters, as in Yaakoubi et al. (2020), (2) a standard CNN-CRF (with non-exhaustive hyperparameter tuning), and (3) our proposed predictor Struct-CKN. We obtain a test error of 0.32%, 0.38%, 0.35% and 0.28% using CNN, CNN-CRF, Struct-CKN-BCFW and Struct-CKN-SDCA, respectively. Therefore, our predictor outperforms both CNN (using hundreds of iterations of GP) and CNN-CRF. Note that although one can use GP to fine-tune hyperparameters, this is not feasible in a real-case usage scenario as practitioners cannot perform GP, each time new data is available.

Computational results on the feasibility and characteristics of initial solutions are summarized in Table 7. We report the number of covered flights, pairings, and the cost of the monthly solution (i.e. cost of undercovered flights + cost of overcovered flights + solution cost of the covered flights). Note that the cost of the monthly solution is inflated and does not represent the real cost of the solution because the cost of undercovered and overcovered flights is expressed in large values. First, note that once we break all illegal pairings in “*Standard - initial*”, 51% of the pairings and 45% of the covered flights are removed, while we only remove 11% of the pairings and 6% of the covered flights for “*Struct-CKN - initial*”.

Table 8 reports the mean of computational results per window for all algorithms, namely, GENCOL-DCA, CNN, Struct-CKN and Struct-CKN+. For each window and each algorithm, we provide the LP value at the root node of the search tree N0 (LP-N0), computational time at N0 (N0 time), the number of fractional variables (# VF-N0) in the current MP solution at N0, number of branching nodes resolved (# Nodes), best LP value found (Best-LP), pairing cost of the best feasible solution (INT) and total computational time (T time); all times are in seconds.

While CNN gives better LP-N0, Best-LP, and INT values with an average reduction factor of 14.37%, 14.17% and 11.48% respectively, compared to GENCOL-DCA, Struct-CKN yields better results with a reduction factor of 32.05%, 32% and 30.94% respectively. On the other hand, note that whilst CNN has a computational time on average 108% larger than that of GENCOL-DCA, Struct-CKN has a computational time only 40% larger on average. This time increase is due to the number of fractional variables at N0 with an increase factor of 81% for CNN and 18% for Struct-CKN, compared to GENCOL-DCA. When base constraints are restrictive, the root node solutions contain a more significant number of fractional-valued pairing variables in order to split the worked time between the bases evenly. This causes an increase in the number of branching nodes required to obtain a good integer solution. While the time increase in CNN is due to the large number of fractional variables at N0, caused by inconsistencies between the initial solution and initial clusters, Struct-CKN has a lower number of nodes and computational time than CNN for two reasons. First, the prediction process is not a greedy process, in the sense that the links (connections) between flights are predicted jointly. This makes the solution more feasible, as shown in Section 6.2.4. Second, and unlike NN-based heuristics, we use the proposed pairings not only as initial clusters, but also as an initial solution. Next, we compare GENCOL-DCA and Struct-CKN+. Struct-CKN+ gives better LP-N0 and Best-LP values for all windows providing an average reduction factor of 56.31% and 55.73% respectively. Likewise, Struct-CKN+ gives better feasible solutions with a reduction factor of 56.26%. This is explained by the ability of Struct-CKN+ to propose a better initial solution, better than those used as training set for the Struct-CKN predictor.

Computational results for the monthly solution obtained at the end of the optimization process are reported in Table 4. We report the total solution cost, cost of global constraints, and number of deadheads obtained with Commercial-GENCOL-DCA fed by GENCOL init (baseline), CNN, CNN+, Struct-CKN, and Struct-CKN+.

While CNN reduces the solution cost and the cost of global constraints by 8.52% and 78.11%, Struct-CKN outperforms CNN reducing the solution cost and the cost of global constraints by 9.51% and 80.25%. Furthermore, while Struct-CKN reduces the number of deadheads in the solution by 7.76%, note that for CNN, the number of deadheads in the solution is slightly larger with an increase factor of 2.17%, compared to GENCOL-DCA. The cost of deadheads is accounted for in the solution cost. Because the solution cost is a multi-objective function, We believe that using slightly more deadheads permitted to get better solutions, enhancing both the solution cost and the cost of global constraints. The solutions found by Struct-CKN+ present better statistics than GENCOL-DCA, CNN and Struct-CKN, with a reduction factor in solution cost and cost of global constraints of 16.93% and 97.24%, respectively. Even more interesting, the number of deadheads used is reduced by 41.23%, compared to GENCOL-DCA, which shows that the Struct-CKN monthly solution can further be optimized and that further research to avoid the windowing approach and use a one-month window can present better results than the current version of Commercial-GENCOL-DCA.

Table 6. Test Error on the Flight-connection Dataset

	Test error (%)	# parameters
CNN (Yaakoubi et al., 2020)	0.32	459 542
CNN-CRF (Chu et al., 2016)	0.38	547 542
Struct-CKN-BCFW	0.35	<b>15 200</b>
Struct-CKN-SDCA	<b>0.28</b>	<b>15 200</b>

**Structured Convolutional Kernel Networks for Airline Crew Scheduling**

*Table 7. Characteristics of Monthly Solutions*

	Flight time	#undercovered flights	#overcovered flights	#covered flights	#pairings	Cost (under. + over. + sol. )	#deadheads
Standard - initial	81804h36	4 960	3 114	45 929	6 525	37 927 328 596	475
Standard - feasible	45644h21	25 838	24	25 051	3 226	29 753 057 875	92
GENCOL init.	82905h03	4 304	23	46 585	6 951	9 199 014 453	1 718
Struct-CKN - initial - extended period	81503h33	5 100	614	45 789	4 916	15 893 302 291	266
Struct-CKN - initial - bid period	75292h22	8 535	599	42 354	4 515	20 294 606 643	241
Struct-CKN - feasible	70672h27	11 148	<b>0</b>	39 741	<b>4 015</b>	16 45 833 4260	220
Struct-CKN + GENCOL	82886h49	4 313	23	46 576	<b>5 993</b>	<b>9 198 925 708</b>	<b>1360</b>

*Table 8. Computational Results per Window*

Win.	Alg.	LP-N0	Diff. (%)	#FV-N0	#Nodes	Best-LP	Diff. (%)	INT	Diff. (%)	T time (s)
1	GENCOL-DCA	10 122 035		2719	159	10 025 487.73		10 276 344.14		<b>7 891</b>
	CNN	9 904 828	-2.15	2870	203	9 891 390.17	-1.34	10 136 106.73	-1.36	10 691
	Struct-CKN	7 099 264	-29.86	2929	200	7 094 683.52	-29.23	7 298 476.40	-28.98	14 855
	Struct-CKN+	<b>2 432 353</b>	<b>-75.97</b>	<b>2413</b>	<b>170</b>	<b>2 423 143.75</b>	<b>-75.83</b>	<b>2 501 388.70</b>	<b>-75.66</b>	14 503
2	GENCOL-DCA	11 396 498		2519	162	11 225 022.17		11 501 016.42		<b>27 778</b>
	CNN	10 319 719	-9.45	5848	419	10 215 252.88	-9.00	10 865 255.80	-5.53	94 004
	Struct-CKN	7 547 443	-33.77	3134	340	7 430 486.16	-33.80	7 915 983.69	-31.17	41 890
	Struct-CKN+	<b>4 533 093</b>	<b>-60.22</b>	<b>2011</b>	<b>113</b>	<b>4 518 960.16</b>	<b>-59.74</b>	<b>4 595 477.05</b>	<b>-60.04</b>	28 210
3	GENCOL-DCA	10 740 127		2330	177	10 641 496.00		11 107 990.12		30 421
	CNN	9 596 326	-10.65	4838	294	9 432 461.69	-11.36	10 051 850.40	-9.51	53 736
	Struct-CKN	7 129 557	-33.62	3045	198	6 977 006.00	-34.44	7 400 639.68	-33.38	34 255
	Struct-CKN+	<b>4 588 242</b>	<b>-57.28</b>	<b>1876</b>	<b>83</b>	<b>4 571 315.93</b>	<b>-57.04</b>	<b>4 650 601.49</b>	<b>-58.13</b>	<b>27 003</b>
4	GENCOL-DCA	9 764 727		2606	229	9 591 592.13		9 968 081.73		33 898
	CNN	8 063 084	-17.43	4763	394	7 868 177.74	-17.97	8 791 799.94	-11.80	56 291
	Struct-CKN	6 427 422	-34.18	3387	183	6 291 641.09	-34.40	6 439 030.07	-35.40	45 430
	Struct-CKN+	<b>4 617 230</b>	<b>-52.72</b>	<b>1938</b>	<b>134</b>	<b>4 601 886.81</b>	<b>-52.02</b>	<b>4 697 567.30</b>	<b>-52.87</b>	<b>27 737</b>
5	GENCOL-DCA	8 095 063		3150	188	7 899 149.01		8 102 703.93		35 948
	CNN	6 076 461	-24.94	5333	417	5 953 932.68	-24.63	6 453 605.80	-20.35	74 622
	Struct-CKN	5 383 225	-33.50	3038	227	5 280 707.35	-33.15	5 560 969.73	-31.37	44 548
	Struct-CKN+	<b>4 331 589</b>	<b>-46.49</b>	<b>1803</b>	<b>89</b>	<b>4 323 521.96</b>	<b>-45.27</b>	<b>4 386 135.29</b>	<b>-45.87</b>	<b>25 196</b>
6	GENCOL-DCA	6 778 925		2513	187	6 610 562.32		6 939 096.41		29 368
	CNN	4 763 520	-29.73	4940	318	4 697 990.54	-28.93	4 947 363.87	-28.70	55 263
	Struct-CKN	5 074 211	-25.15	3222	278	5 003 421.94	-24.31	5 368 766.13	-22.63	51 166
	Struct-CKN+	<b>4 356 910</b>	<b>-35.73</b>	<b>1961</b>	<b>126</b>	<b>4 349 221.21</b>	<b>-34.21</b>	<b>4 491 319.98</b>	<b>-35.28</b>	<b>26 529</b>
Mean	GENCOL-DCA	9 482 896		2640	184	9 332 218.23		9 649 205.46		27 551
	CNN	8 120 656	-14.37	4765	341	8 009 867.62	-14.17	8 540 997.09	-11.48	57 435
	Struct-CKN	6 443 520	-32.05	3126	238	6 346 324.34	-32.00	6 663 977.62	-30.94	38 691
	Struct-CKN+	<b>4 143 236</b>	<b>-56.31</b>	<b>2000</b>	<b>119</b>	<b>4 131 341.64</b>	<b>-55.73</b>	<b>4 220 414.97</b>	<b>-56.26</b>	<b>24 863</b>

*Table 9. Computational Results for Monthly Solutions*

Initial solution and clusters	Solution cost	Diff. vs GENCOL-DCA (%)	Cost of global constraints	Diff. vs GENCOL-DCA (%)	Number of deadheads
GENCOL init	30 681 120.5		9 465 982.28		1725
GENCOL-DCA	20 639 814.6	-32.73 (vs GENCOL init)	2 127 086.77	-77.53 (vs GENCOL init)	992
CNN	18 881 977.9	-8.52	465 687.94	-78.11	1014
Struct-CKN	<b>18 677 641.1</b>	<b>-9.51</b>	<b>420 095.87</b>	<b>-80.25</b>	<b>915</b>
CNN+	18 617 109.4	-9.80	333 983.69	-84.30	997
Struct-CKN+	<b>17 145 543.4</b>	<b>-16.93</b>	<b>58 797.21</b>	<b>-97.24</b>	<b>583</b>