Les Cahiers du GERAD ISSN: 0711–2440

Comptes rendus du 12e atelier de résolution de problèmes industriels de Montréal, 22–26 août 2022

Proceedings of the 12th Montréal Industrial Problem Solving Workshop, August 22–26, 2022

Odile Marcotte, éditrice

G–2022–63

Décembre 2022

Citation suggérée : Odile Marcotte, éditrice (Décembre 2022). textcolorredComptes rendus du 12e atelier de résolution de problèmes industriels de Montréal, 22–26 août 2022 / Proceedings of the 12th Montréal Industrial Problem Solving Workshop, August 22-26, 2022, Rapport technique, Les Cahiers du GERAD G–2022–63, GERAD, HEC Montréal, Canada.

Suggested citation: Odile Marcotte, éditrice (December 2022). textcolorredComptes rendus du 12e atelier de résolution de problèmes industriels de Montréal, 22–26 août 2022 / Proceedings of the 12th Montréal Industrial Problem Solving Workshop, August 22-26, 2022, Technical report, Les Cahiers du GERAD G–2022–63, GERAD, HEC Montréal, Canada.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2022
– Bibliothèque et Archives Canada, 2022

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2022
– Library and Archives Canada, 2022

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Préface

Le Douzième atelier de résolution de problèmes industriels de Montréal, qui eut lieu du 22 au 26 août 2022, fut organisé conjointement par le Centre de recherches mathématiques (CRM) et l'Institut de valorisation des données (IVADO). La préparation de l'atelier fut marquée par le retour au mode habituel, où les participants se rencontrent en personne, après deux années d'ateliers virtuels. Plus de 80 personnes s'inscrivirent à l'atelier et examinèrent huit problèmes, fournis respectivement par Air Canada, la Banque Nationale du Canada, Beneva, Environnement et Changement climatique Canada, Hydro-Québec, l'IATA, Radio-Canada et Revenu Québec. Je remercie chaleureusement ces partenaires et les coordonnateurs des équipes (Frédéric Quesnel, Gilles Caporossi, Juliana Schulz, Julie Carreau, Alexandre Blondin Massé, Mike Lindstrom, Marie-Jean Meurs et Alain Hertz), ainsi que les conseillers IVADO dont la collaboration fut essentielle à la collecte des problèmes. Finalement j'exprime toute ma reconnaissance à Karine Hébert, qui m'a aidée à mettre en forme ces comptes rendus.

Odile Marcotte
Conseillère spéciale aux partenariats, CRM
Professeure associée, UQAM et membre associé, GERAD

CENTRE
DE RECHERCHES
MATHÉMATIQUES

IVADO

Foreword

The Twelfth Montreal IPSW took place on August 22-26, 2022, and was jointly organized by the Centre de recherches mathématiques (CRM) and the Institute for Data Valorization (IVADO). For the first time since 2019, the workshop was held "in person," after two virtual workshops held in 2020 and 2021, respectively. More than 80 persons registered for the workshop and studied eight problems, submitted by Air Canada, the National Bank of Canada, Beneva, Environment and Climate Change Canada, Hydro-Québec, IATA, Radio-Canada, and Revenu Québec. I am grateful to our industrial partners and the team coordinators (Frédéric Quesnel, Gilles Caporossi, Juliana Schulz, Julie Carreau, Alexandre Blondin Massé, Mike Lindstrom, Marie-Jean Meurs, and Alain Hertz), as well as the IVADO advisors who helped us find the problems submitted to the workshop. I am also very grateful to Karine Hébert, who helped me put these proceedings together.

Odile Marcotte
Special Advisor, Partnerships, CRM
Adjunct Professor, UQAM and Associate Member, GERAD

CENTRE
DE RECHERCHES
MATHÉMATIQUES

IVADO

Contents

1 Building a maintenance plan

Mohammad Rahim Akhavan [c]
Warley Almeida Silva [a]
Samaha Abdalrhaman [b]
Keith Dugas [c]
Prakash Gawas [d]
Hewan Leul [d]
Frédéric Quesnel [b]

[a] *DIRO, Université de Montréal*
[b] *UQAM*
[c] *Air Canada*
[d] *Polytechnique Montréal*

December 2022
Les Cahiers du GERAD
Copyright © 2022 GERAD, Akhavan, Silva, Abdalrhaman, Dugas, Gawas, Leul, Quesnel

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

1.1 Introduction

Air Canada has a large fleet of heterogeneous aircraft and maintenance on those aircraft is a complex logistical problem. Maintenance has to be carried out during ground time (i.e., time spent on the ground between flights) and according to the resources available at the location of the aircraft (i.e., airport).

At the moment maintenance schedules are prepared manually and collaboratively by around 30 planners. Air Canada wishes to implement an automated system instead. Furthermore the company wishes to integrate two types of maintenance tasks: scheduled maintenance and defect repairs.

In this project we were asked to design a new way of creating a maintenance schedule for all aircraft in the Air Canada fleet. The schedule should be feasible for a long-term horizon (2 years) and optimized in the short term (around 30 days). It should also be possible to create a full schedule relatively quickly (in minutes). This would allow the company to generate a new schedule once a batch of new information arrives (e.g., a list of new defects).

For a maintenance task to be carried out on the plane enough manpower must be available. We are given the daily availability in terms of head counts and we first decide how to allocate the manpower to shifts for the day over the planning period. This will directly affect when maintenance tasks are scheduled. The goal is also to optimize indirectly the utilization of manpower. We use a heuristic that allocates manpower in a proportional fashion according to the percentage of flights arriving during those shifts. The task assignment heuristic sorts tasks according to a prioritization score that takes into consideration the number of man hours and the task type, and progressively schedules tasks to the earliest (resp. farthest) available ground time if it is a defect repair (resp. scheduled maintenance), which implicitly minimizes (resp. maximizes) yield remaining.

The remainder of this report is structured as follows. In Section 1.2 we describe the maintenance schedule problem. The available data is presented in Section 1.3. Sections 1.4 and 1.5 present the two approaches that were considered to solve the maintenance schedule problem. The first one is based on mathematical programming and the second one is a greedy algorithm. Results obtained with the greedy approach are presented in Section 1.6 and concluding remarks are found in Section **??**.

1.2 Problem description

This section describes the maintenance schedule problem proposed by Air Canada. The goal of this problem is to create a feasible maintenance schedule while optimizing a multi-objective function. We first present and motivate this objective function in Section 1.2.1. The requirements for feasible maintenance schedules are given in 1.2.2. A preliminary step of the maintenance schedule problem is assigning ground crews to work shifts. Details on this step are presented in Section 1.2.3. Section 1.2.4 includes a discussion on the separability of the scheduling problem by fleet type. Finally Section 1.2.5 includes additional considerations that should be taken into account in a real-life implementation of our method but were ignored during the workshop because of the limited time available.

1.2.1 Objective

Optimal scheduling decisions depend on whether the maintenance task is a *scheduled maintenance task* or a *defect repair task*. Scheduled maintenance is maintenance that has to be performed at specific intervals, such as changing tires, oiling the engine, and replacing parts that deteriorate over time. The intervals at which maintenance tasks have to be performed depend on several factors, such as the number of flights and the elapsed time, but they can be translated into a due date. It is desirable to carry out scheduled maintenance as close as possible to the due date in order to utilize airplane parts as much as possible. This is referred to as *minimizing the yield remaining.*

Each defect repair task has a *priority code* indicating its priority level. A good maintenance schedule will perform those repair tasks first. Defect repair tasks concern fixing non-critical issues in the airplane that do not prevent it from operating safely (e.g., entertainment system malfunction, broken seat, etc). Although defect repair tasks are non-critical, they also have assigned due dates based on company standards. It is preferable to correct the issue as soon as possible in order to minimize a disruption leading to a passenger's discomfort. In other words defect repair tasks should be carried out as far as possible from the due date. This is referred to as *maximizing the yield remaining*.

1.2.2 Feasible maintenance schedule

Although the flight schedule is known well in advance, the aircraft routes are typically known only a few days in advance. This is somewhat problematic because the flight schedule only indicates when and where aircraft of some fleet type will be available for maintenance, but not which specific aircraft will be available. To address this issue we assume that we can choose which aircraft of some fleet type will be available during each ground time for that fleet type. The schedule will be adapted as aircraft routes become available.

A time interval during which maintenance can be performed is called a *ground time*. A ground time is defined by a *location*, an *fleet type*, and a *time interval*.

Furthermore each location has a limited number of ground crews available to perform repairs. Each crew is qualified to perform maintenance on a specific fleet type. Each crew works on a *cycle* and a *shift*. The cycle determines the work dates and the shift determines the work hours.

To perform a maintenance task on an aircraft in any compatible ground time, the following conditions have to be satisfied.

- The ground time has to be long enough for the task to be carried out.
- Enough qualified crews must be available to work on the task. The number of crews assigned per task is variable and the task completion time depends linearly upon that number. For instance a four-hour task can be completed in two hours by two crews or in one hour by four crews.

It is possible to complete several tasks in any ground time. All tasks carried out in the ground time, however, must be related to the same aircraft. This is because ground time represents the time when a specific aircraft is available (even though we don't know which one).

1.2.3 Assigning crews to shifts

Each crew is associated with a given location, fleet type qualification, *cycle*, and *shift*. A crew can only perform work at their designated location. Similarly a crew can only work on aircrafts corresponding to their qualification. A cycle defines on which days a crew will work and follows a regular pattern. A shift corresponds to the working hours of each crew. We assume all crew keep the same cycle and shift throughout the optimization period. Although the cycle of each crew is considered fixed in the input, their shifts are not. Judiciously selecting the crew shifts can have a large impact on the quality of the maintenance schedule, because it determines the workforce available when maintenance is performed on an aircraft. Ideally more crews would be available in periods where ground times occur.

1.2.4 Separability by fleet type

Even though ground times do not have a preassigned airplane, they require an airplane of some fleet type. At each location, each shift and cycle has a number of employees able to work on a preassigned fleet type. Note that there is no resource shared between fleet types, and that each fleet type has its own subset of tasks to assign to its own subset of ground times. In this sense our scheduling problem can be separated by fleet types. This means that a schedule built by merging optimal schedules for their respective flight types is optimal for the overall problem. Our approaches will make use of this property: they solve smaller problems and merge their optimal solutions in order to obtain a global one.

1.2.5 Simplifications

Because of the limited duration of the workshop, simplifications had to be made to the "real" maintenance schedule problem. We list some of them below.

Hangar constraints: While shorter maintenance tasks can be performed on the tarmac, longer maintenance tasks require bringing the aircraft to a hangar. This substantially complicates the problem because it requires taking into account the travel time. Furthermore this version of the problem would no longer be separable by fleet type, as each hangar has a limited capacity.

Maximum/minimum number of crews per task: In real life it is not always possible to add more crews to perform a task faster. Likewise some tasks require more than one crew to complete (if the task requires lifting heavy parts, for instance). We ignored those restrictions, however, because the available data did not specify a maximum or a minimum number of crews per task.

Conflicting tasks: Our model supposes that any two tasks can be performed at the same time. This is not the case in practice, however, as some tasks are mutually conflicting. We ignored those restrictions because the data did not specify which tasks are incompatible.

Number of ground crews per fleet type: Ground crews can be retrained to work on a different fleet type. It would be interesting to do some analysis to determine which fleet types don't have a sufficient number of crews.

1.3 Available data

The report is based entirely on data obtained for Air Canada by Optimus|AI in 2022. In this section we will be looking at the data and the methodology we use to analyze them.

1.3.1 Description of the data

We have seven collections of data: Defects, MPD, Ground Times, Headcount by Cycle, HeadCount, Shift Start Times, and Schedule Cycles. Here is their description.

Defects: FIN Type, Fleet Type, DEFECT_TYPE (L, M, C, E), DEFECT_NUMBER, DEFECT_ITEM, PLAN_PROJECTED_DUE_DATE, MAN_HOURS_VALUE, Due_Date, and Due_time.

MPD: FIN Number, Fleet Type (F-192, F-231, F-295, F-546, F-630, F-670, F-690, F-855), EO, TASK_CARD, DUE_DATE, MAN_HOURS_VALUE, Due_Date, and Due_time.

Ground Times: Flight Local Date, Fleet Type, Departing Station, Inbound Arrival Local Time, Ground Minutes, Arriving Station, Departure Local Time, and Outbound Arrival Local Time.

Headcount by Cycle: Cycle 1 (Location, FLEET), Cycle 2 (Location, FLEET) for each month from July 2022 to January 2023.

HeadCount: Month, Cycle, Location (YEG, YHZ, YOW, YQB, YUL, YVR, YWG, YYC, YYT, YYZ), and F670|F192|F546, F295, F690, F855, F630, F231, F286.

Shift Start Times: Four shift starts from (06:30 – 17:40, 11:00 – 22:10, 16:00 – 03:10, 20:00 – 07:25).

Schedule Cycles: Line Maintenance with 1 / 2 Rotation 2022 for each month from July 2022 to January 2023.

1.4 First approach : Integer programming

This section presents the integer program initially proposed for modelling the problem.

1.4.1 Input information

Given that the scheduling problem is separable by fleet type, we propose a model for a predetermined fleet type. In what follows $\mathcal{G}$ denotes the set of ground times, indexed by g; $\mathcal{L}$ the set of locations (airports), indexed by ℓ; $\mathcal{D}$ the set of defect repairs, indexed by d; $\mathcal{M}$ the set of maintenance tasks, indexed by m; $\mathcal{T}$ the set of tasks ($\mathcal{T} = \mathcal{D} \cup \mathcal{M}$), indexed by t; $\mathcal{C}$ the set of cycles ($\mathcal{C} = \{1, 2\}$), indexed by c; and $\mathcal{S}$ the set of shifts ($\mathcal{S} = \{1, 2, 3, 4\}$), indexed by s.

For some ground time g, let $\mathcal{C}(g)$ be the subset of cycles that overlap with the ground time g (i.e., cycles that happen on the same day(s) as the ground time g); $\mathcal{S}(g)$ be the subset of shifts that overlap with the ground time g (i.e., shifts that happen in the same hour(s) as the ground time g); and $\mathcal{L}(g)$ be the subset of locations that overlap with the ground time g (i.e., the locations where the ground time g happens).

Also let $c_t^g \in \mathbb{Z}$ be the penalty from letting task t being performed at ground time g; $h_{cs}^g \in \mathbb{Z}$ be the number of workforce hours that ground time g has on shift s and cycle c; $b_c^{\ell-}$ be the lower bound on the number of people that can work on cycle c; and $b_c^{\ell+}$ be the upper bound on the number of people that can work on cycle c.

These model parameters can be easily computed with the available data. More specifically, the penalty parameters c_t^g can be computed as the distance in the number of days between the planning day and the ground time g for defect repair tasks (i.e., larger penalties would mean that defects have remained unfixed for longer), and the distance in the number of days between the ground time g and the due time for task t for scheduled maintenance tasks (i.e., larger penalties would mean that maintenance has been carried out sooner than needed).

1.4.2 Mathematical model

The integer program has the following decision variables. Binary variable x_t^g equals 1 if ground time g is assigned task t, 0 otherwise. Integer variables y_{cs}^ℓ describe the number of people on shift s in cycle c at location ℓ. Integer variables $w_g \in \mathbb{Z}$ describe the number of people available to work during ground time g.

Given these decision variables, the integer program can be written as follows.

$$\min_{w,x,y} \quad \sum_{g \in \mathcal{G}} \sum_{t \in \mathcal{T}} c_t^g \cdot x_t^g \tag{1.1}$$

$$\text{s.t.:} \quad \sum_{g \in \mathcal{G}} x_t^g = 1 \qquad \forall t \in \mathcal{T} \tag{1.2}$$

$$\sum_{t \in \mathcal{T}} d_t \cdot x_t^g \leq w_g \qquad \forall g \in \mathcal{G} \tag{1.3}$$

$$w_g = \sum_{c \in \mathcal{C}(g)} \sum_{s \in \mathcal{S}(g)} h_{cs}^g \cdot y_{cs}^{\mathcal{L}(g)} \qquad \forall g \in \mathcal{G} \tag{1.4}$$

$$b_c^{\ell-} \leq \sum_{s \in \mathcal{S}} y_{cs}^\ell \qquad \forall c \in \mathcal{C}, \forall \ell \in \mathcal{L} \tag{1.5}$$

$$\sum_{s \in \mathcal{S}} y_{cs}^\ell \leq b_c^{\ell+} \qquad \forall c \in \mathcal{C}, \forall \ell \in \mathcal{L} \tag{1.6}$$

$$x_t^g \in \{0, 1\} \qquad \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \tag{1.7}$$

$$y_{cs}^\ell \in \mathbb{Z} \qquad \forall c \in \mathcal{C}, \forall s \in \mathcal{S} \tag{1.8}$$

$$w_g \in \mathbb{Z} \qquad \forall g \in \mathcal{G} \tag{1.9}$$

The objective function (1.1) minimizes the penalties from assigning tasks to ground times. Note that penalties are computed based on the goal of minimizing or maximizing the yield remaining. Constraints (1.2) guarantee that each task must be assigned to a ground time. Constraints (1.3) guarantee that the duration of tasks assigned to some ground time respects the workforce available during that ground time. Constraints (1.4) establish the workforce available at each ground time based on how shifts and cycles have been set by the respective decision variables. Constraints (1.5) and (1.6) guarantee that the number of people assigned to some shift and some cycle respect the lower and upper bounds established by the company. Constraints (1.7)–(1.9) establish the previously defined domains of the decision variables.

1.4.3 Challenges and shortcomings

The integer program initially proposed for the problem could not take into consideration two key characteristics of this problem.

First the integer program assumes that each ground time has a fixed airplane, which is not the case as mentioned previously. Instead of having a fixed airplane, each ground time has a fixed fleet type: the scheduling approach should assign an airplane to some ground time on-the-fly. The model could be adapted as well to assign available airplanes to ground times based on the needed fleet type by adding an airplane index to the variables x_t^g, rewriting the existing constraints appropriately, and adding constraints enforcing valid assignments. Such an extension, however, would increase the number of binary variables in the model, especially if there is a large number of airplanes for some fleet type.

Second the integer program does not properly handle cases where the workforce has been split between two or more ground times. For example, if there were 10 man-hours scheduled during ground times 1 and 2 overlapping completely (i.e., beginning and ending at the same times), the integer program would have as a feasible solution assigning up to 10 man-hours to each ground time, thus not verifying that such a solution would require up to 20 man-hours. This is not a reasonable assumption in practice, as the available workforce is not able to work on two different planes at the same time and should be split accordingly. Adding such a constraint to the model would require some discretization of time, which would not be practical for such a long planning horizon.

One way to overcome this difficulty would be to forbid the workforce to work on two or more ground times that overlap. Then there would be a constraint for each pair of ground times that overlap with each other. Another solution would be to have an external loop solving the integer program at each iteration and adding a constraint blocking infeasible allocations of the workforce once they appear. In other words the algorithm would progressively block solutions that do not respect the splitting of the workforce. Unfortunately such an approach would probably take a long time to find an optimal solution because there seems to be many overlapping pairs of ground times. In the worst-case scenario there would be an iteration for each overlapping of ground times and the algorithm would have to solve multiple integer programs.

1.5 Second approach : Greedy heuristics

This section presents the greedy heuristic proposed for the problem.

1.5.1 Structure overview

Given that the scheduling problem is separable by fleet type, in the sequel we consider that the problem will be solved for a predetermined fleet type. Since solving the problem by mathematical modeling is a challenge, our team decided to propose a heuristic approach and break the problem into two parts.

The first part, referred to as *shift profiles*, aims to identify what is the most reasonable way to build the workforce schedule (i.e., how to assign workers to shifts and cycles). To this end our team proposes a data-driven approach that assigns more workers to time windows with more ground times.

The second part, referred to as *task assignment*, aims to identify what is the best assignment of tasks to ground times. Implicitly, this part also assigns specific airplanes to ground times on-the-fly, as airplanes routes are not known in advance. To this end, our team proposes a greedy approach that assigns tasks to ground times based on a prioritization score. Tasks with a higher score are assigned first and they are to be carried out as early as possible if they are defect repair tasks or as late as possible if they are maintenance tasks.

Sections 1.5.2 and 1.5.3 present the first and second parts of our approach, respectively.

1.5.2 Shift profiles

The first part of the solution process is to assign employees to shifts. Any location has a fixed capacity of contracted employees that work in shifts. The shifts can be classified as follows:

- Morning(1)
- Afternoon(2)
- Evening(3)
- Night(4)

Each shift is approximately 11 hours long. For the sake of simplicity, however, we only consider a six-hour time window called *quarter* for each shift. A simple way of distributing the available manpower would be to distribute it equally among the shifts. This may not be the right thing to do, however, since the arrival pattern at any given airport may not be uniform. We analyze the arrival patterns for each airport. In the table, the available manpower at each airport is given. We see most manpower is distributed at YYZ, YUL, YVR.

N_Fleet_type	**YEG**	**YHZ**	**YOW**	**YQB**	**YUL**	**YVR**	**YWG**	**YYC**	**YYT**	**YYZ**
F-231	0	0	0	0	10	10	0	3	0	32
F-286	1	1	1	1	16	14	0	3	1	35
F-295	2	1	1	0	19	18	1	2	0	50
F-347	2	2	2	1	40	36	2	6	2	92
F-630	0	0	0	0	22	17	0	4	0	49
F-690	0	0	0	0	4	4	0	0	0	15
F-855	0	0	0	0	14	11	0	2	0	36

Given below is a plot of arrivals distribution for each fleet during each quarter.

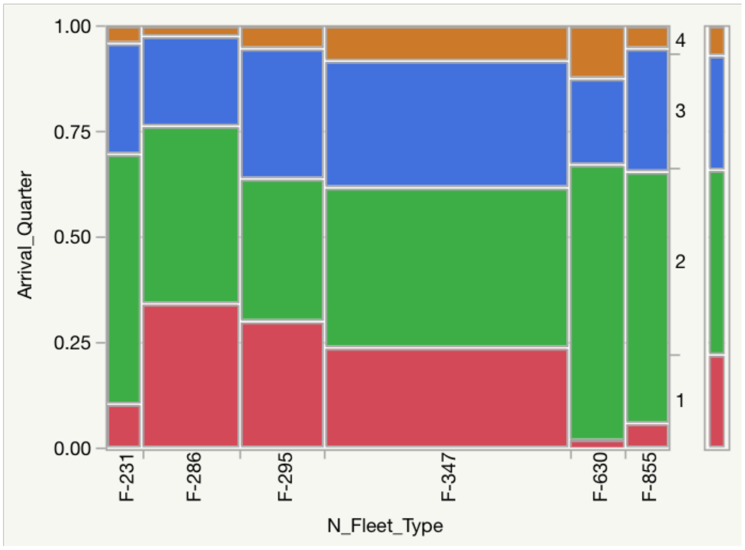

Figure 1.1: Percent of arrivals for a given quarter and fleet type at Toronto.

From Figure 1.1 it is clear that the arrivals are unbalanced during the day. Quarter 4 sees the least arrivals overall in fleets. Quarter 1 has fewer arrivals for certain fleets also. Hence one idea would be to distribute the available manpower according to arrival percentages at each location for each fleet. Suppose X is the total available fleet at a location for a given day. Also let p_q^{fl} be the percent of fleet f arriving during quarter q at location l. Then the number of employees assigned at location l for fleet f

and quarter q would be equal to $n_q^{fl} = X * p_q^{fl}$. In the implementation, we artificially set p_q^{fl} to zero if it is less than 10% percent of the total arrivals at the location and normalize again. The table below gives the distribution of arrivals at the three main airports. The final manpower distribution is given in Figure 1.2.

N_Fleet_Type	Quarter	YUL	YVR	YYZ
F-231	N(1)	0	0.8	0.11
F-231	N(2)	0.81	0	0.62
F-231	N(3)	0.19	0.2	0.27
F-231	N(4)	0	0	0
F-286	N(1)	0.26	0.41	0.35
F-286	N(2)	0.5	0.18	0.43
F-286	N(3)	0.24	0.41	0.22
F-286	N(4)	0	0	0
F-295	N(1)	0.29	0.41	0.32
F-295	N(2)	0.33	0.29	0.36
F-295	N(3)	0.27	0.3	0.32
F-295	N(4)	0.12	0	0
F-347	N(1)	0.36	0.32	0.26
F-347	N(2)	0.3	0.48	0.41
F-347	N(3)	0.34	0.19	0.33
F-347	N(4)	0	0	0
F-630	N(1)	0.46	0.65	0
F-630	N(2)	0.39	0.18	0.67
F-630	N(3)	0.14	0.17	0.21
F-630	N(4)	0	0	0.12
F-855	N(1)	0.18	0.8	0
F-855	N(2)	0.82	0	0.67
F-855	N(3)	0	0.2	0.33
F-855	N(4)	0	0	0

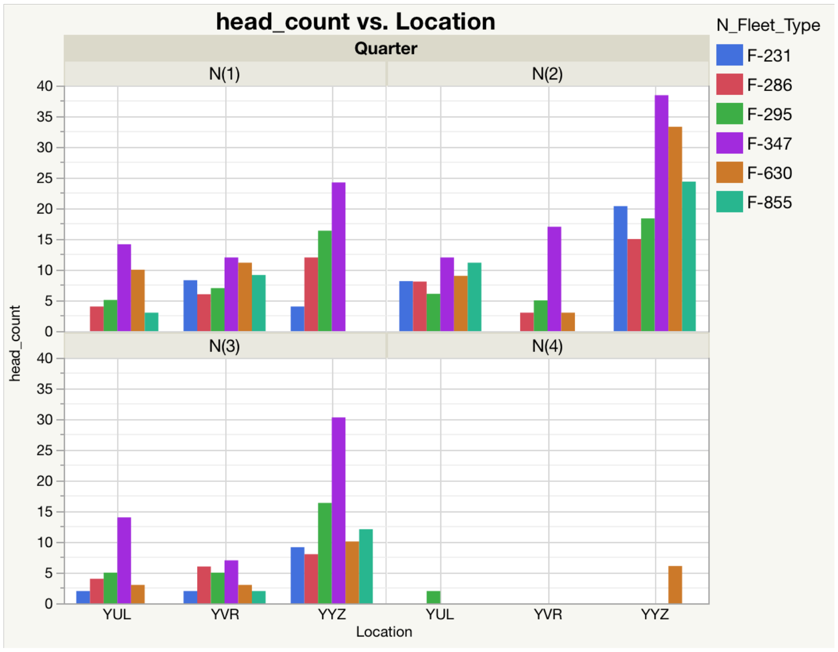

Figure 1.2: Number of employees for each shift and fleet type by location.

1.5.3 Task assignment

The second part of the solution process is to assign tasks to ground times given a shift profile. The greedy algorithm starts by ranking (a) tasks according to a prioritization score in decreasing order (i.e., tasks with a higher score come first) and (b) ground times according to the due dates in increasing order (i.e., ground times closer to the planning day come first). Here is the proposed greedy algorithm.

Step 1: Rank tasks in decreasing order of prioritization score. Go to step 2.

Step 2: Rank ground times from the earliest date to the latest date. Go to step 3.

Step 3: Create empty schedule of tasks for ground times. Go to step 4.

Step 4: Create empty assignment of airplanes to ground times. Go to step 5.

Step 5: Create empty list to store nonassignable tasks. Go to step 6.

Step 6: If there remains a task that has not been analyzed, go to step 7. Otherwise, go to step 12.

Step 7: Choose the task t with the highest prioritization score. Go to step 8.

Step 8: Choose the closest (resp. farthest) ground time g without an assigned airplane or with the same airplane as task t if task t is a defect repair (resp. scheduled maintenance) task. Go to step 9.

Step 9: If there is no ground time g able to host task t, add task t to the list of unassigned tasks and go to step 6. Otherwise, go to step 10.

Step 10: Compute the maximum number of workers w available to work on task t during ground time g based on previous assignments. Go to step 11.

Step 11: Schedule task t to ground time g with w workers. If ground time g has no assigned airplane, assign it to the airplane of task t. Go to step 6.

Step 12: Export the list of unassigned tasks and the schedule to files and stop.

The greedy algorithm stops once all tasks have been analyzed and either scheduled for some ground time with a number of workers or deemed nonassignable due to previous assignments and properties of tasks. Some steps of the greedy algorithm are straightforward (e.g., ranking ground times from the earliest date to the latest date) but others require a careful definition. In particular there are two challenges in the implementation of the greedy approach. The first challenge arises in Step 1: how does one compute the prioritization score of some task t? The second challenge arises in Step 10: how many employees are available during some time period given previously scheduled tasks? We discuss these challenges in detail below.

Prioritization score

The prioritization score takes into consideration two attributes of the task. The first attribute is the task priority. There are four priorities for defect repair tasks: L, M, C, and E. By assigning the letter S for scheduled maintenance tasks, we have five task priorities altogether. In practice the priorities follow these rules: L $\succ$ S $\succ$ M, C $\succ$ E (i.e., L has the highest priority, followed by S; M and C have the same priority; and E has the lowest priority). The second attribute is the number of man-hours required to complete the task. In practice tasks with a large number of man-hours are scheduled first because they require long ground times and a larger number of employees. The prioritization score of a task t can be computed as

$$score(t) = manhours(t) + priority(t),$$

where $manhours(t)$ returns the number of man-hours required for that task and $priority(t)$ returns a number corresponding to the task priority (one of the five types described above). Function $priority(t)$ has been set to 10 for task type L, 8 for task type S, 6 for task types M and C, and 4 for task type E. Intuitively tasks with a large number of man-hours will always have higher prioritization scores, whereas tasks with a lower number of man-hours will be prioritized based on a trade-off between man-hours and task priority. Note that the definition of the function $priority(t)$ could be improved based on a further analysis of the algorithm.

Availability profile

In order to assess whether a given task can be carried out within a given ground time, it is necessary to know whether enough crews are available during that time. This is somewhat challenging because the number of available crews changes over time as work shifts begin and end. Furthermore our greedy algorithm must take into account the tasks already assigned, which lowers the number of crews for some time intervals.

We keep track of the ground crews availability using a *crew availability profile*, which can be visualized as a plot of availability as a function of time. Such a graph is displayed in Figure 1.3.

Since the workforce is split by fleet type and location, we create one crew availability profile for each combination of fleet type and location. Each time a task is assigned to some ground time, the appropriate availability profile is updated to reflect the reduced availability during that time.

To store a shifting profile it is only necessary to keep a list of *events* where the availability changes, as well as the new availability after those events. The red points in Figure 1.3 represent the data points that must be stored.

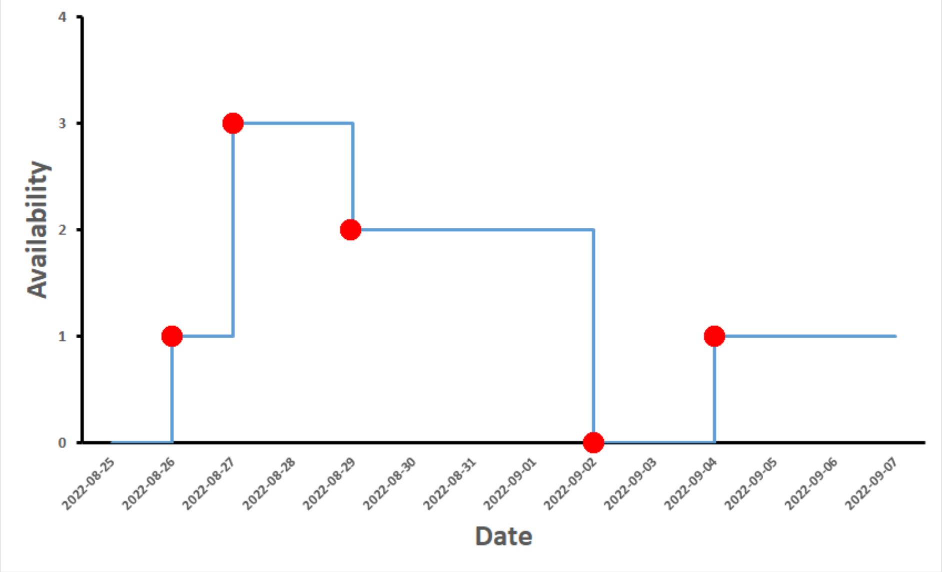

Figure 1.3: Example of shift profile.

1.5.4 Discussion of the greedy algorithm

The greedy algorithm could be improved in several ways. First note that the prioritization score of each task remains constant throughout the algorithm. In practice it could be beneficial to update the prioritization score (and the corresponding ordering of the tasks) as the greedy algorithm progresses. For instance, once a task for airplane a has been assigned to ground time g, it may be helpful to consider in priority the other tasks for a (to see if they fit in ground time g). Similarly one could compute a prioritization score for each task/ground pair. We experimented with such approaches, but they proved too computationally demanding. We believe a better implementation of the algorithm would be significantly faster and would allow for more sophisticated prioritization methods.

Another limitation of our approach is that selecting shift profiles and performing task assignment are done sequentially. A better approach would be to integrate the two steps. We propose two ways to do this. First shift profiles could be determined inside the greedy algorithm. This approach may not be the best because each shift profile decision has a long-term impact that the greedy algorithm may not be able to take fully into account. The second possibility is to iterate between the shift profile step and the task assignment step. Once task assignment is performed, more information is available to create better shift profiles, which will lead to better task assignments. Implementing this approach would require formulating the shift profile step as an optimization problem, which should be relatively easy to do.

1.6 Results

To quantify the quality of our solution we came up with two main metrics. The first is the yield time remaining (YTR), which is the difference in time units between the day the task was carried out and the due date. Note that for a maintenance task the preferred start date is a date that is as late as possible and for a repair task the date needs to be as early as possible. The second metric measures the utilization of the available resources at each airport. In Figure 1.4 we have the distribution of YTR for the maintenance tasks. More than 90% of the tasks are scheduled within 5 days of their due date. There are still some outliers, however, that we need to schedule much earlier than the due date. This may be due to longer task times and the lack of availability of ground time for these tasks. In Figure 1.5 we display the distribution of YTR for the repair tasks. We manage to schedule more than 95% of the tasks within 5 days of their due date.

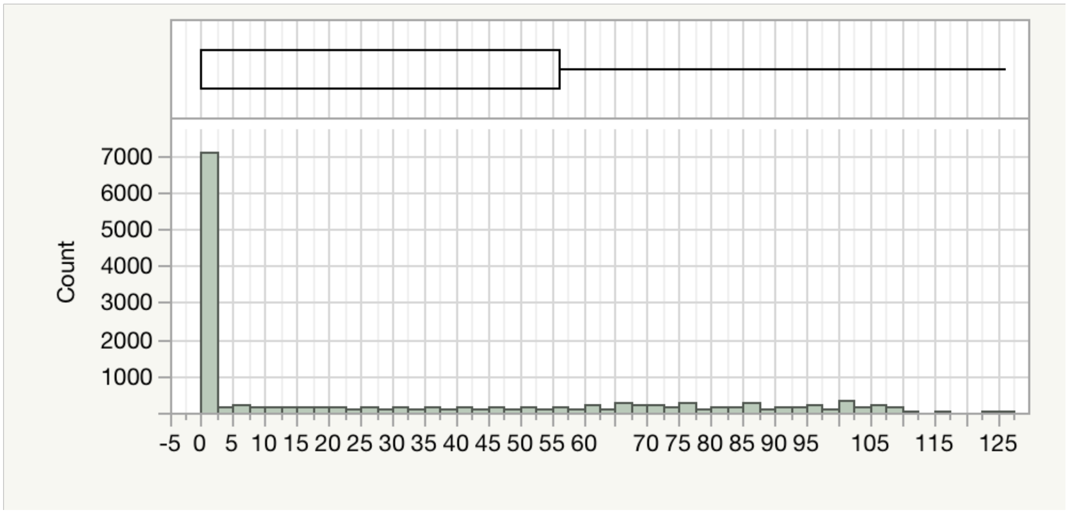

Figure 1.4: Yield remaining for maintenance tasks.

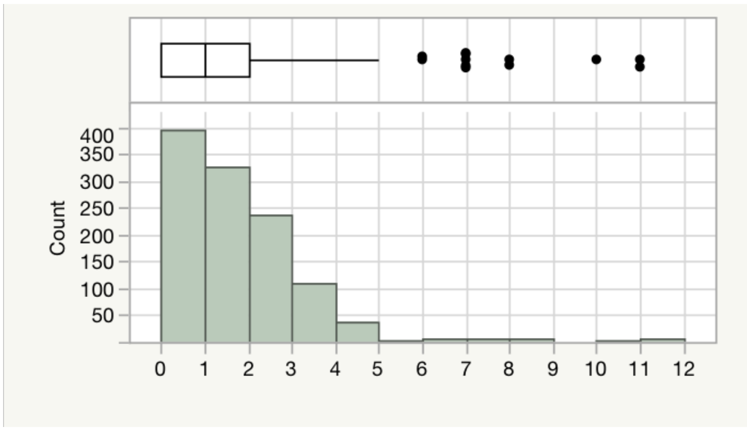

Figure 1.5: Yield remaining for repair tasks.

Figures 1.6 and 1.7 give the fleet-wise break-up of means and 90^{th} quantile of the YTR for repair and maintenance tasks, respectively. We can see that certain fleets have a larger YTR than others. It may be that these fleets have a low availability of ground time.

Figure 1.8 gives the resource utilization for Toronto airport for fleet F-286. We compiled such curves for the main airports, but here we present only one such curve for illustration. As can be seen a lot of manpower seems idle; and as most tasks have the same due date they are bunched together. Based

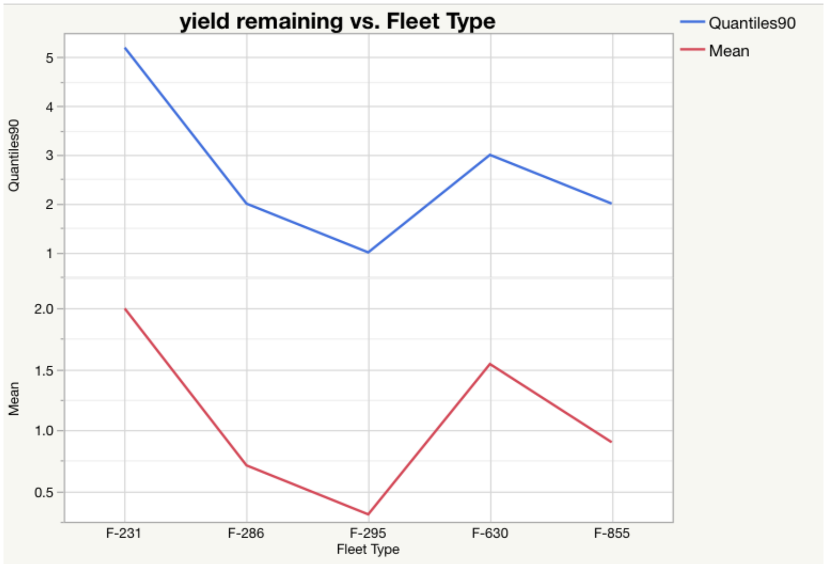

Figure 1.6: Yield remaining for repair tasks fleet-wise.

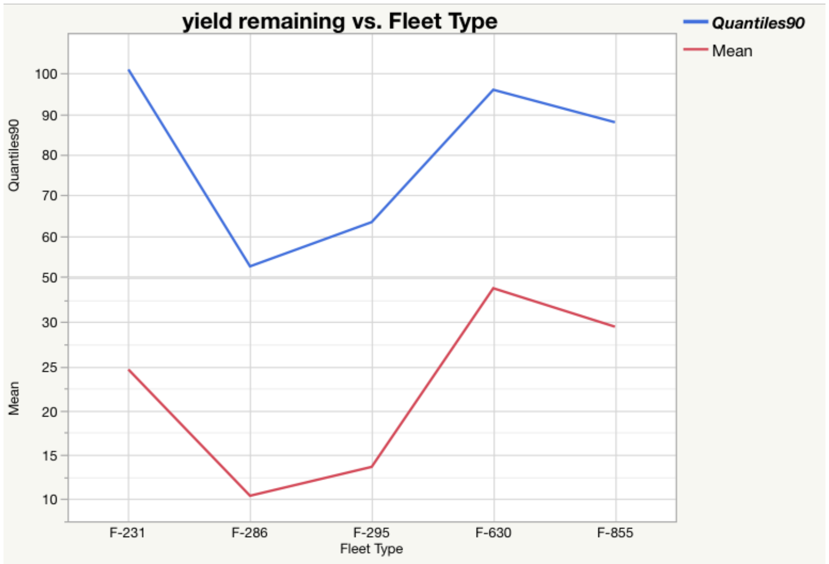

Figure 1.7: Yield remaining for maintenance tasks fleet-wise.

on these observations the company can conduct further research and come up with better workforce policies over the planning period.

Finally the table below gives the runtime consumed by the algorithm for finding the greedy solution.

Fleet type	Runtime (minutes)
F-231	11.2
F-286	15.1
F-295	17.7
F-630	25.4
F-855	8.1

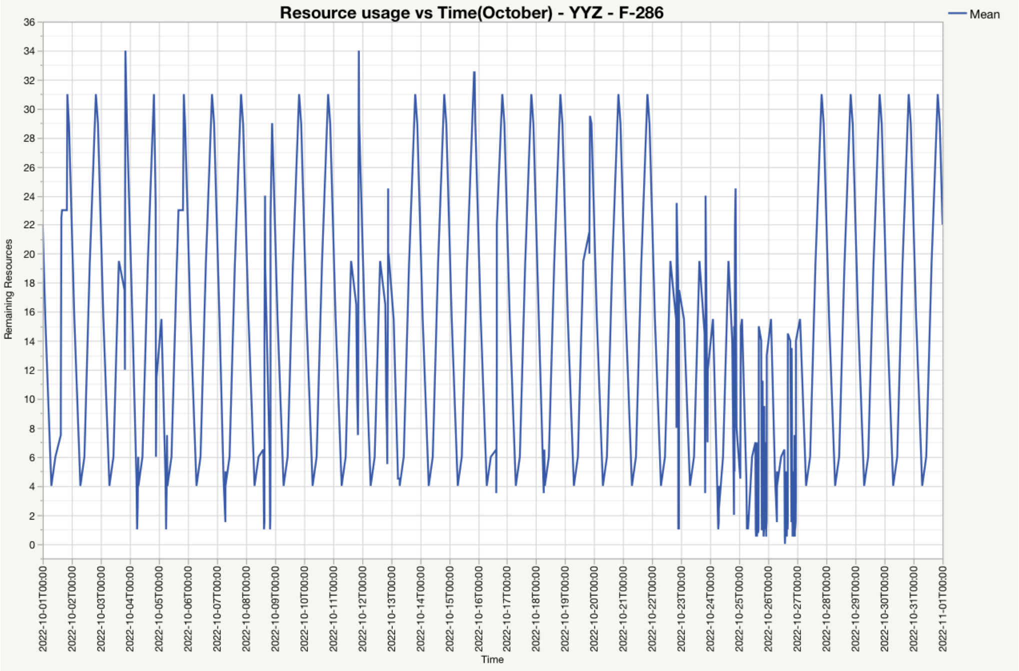

Figure 1.8: Resource utilization for Toronto for fleet F-286.

2 Beneva: Survival models and incomplete data

Golshid Aflaki [a]
Eloi Bizimana [b]
Ariane Cwiling [c]
Faramarz Farhangian [d]
James McVittie [e]
Bruno Monsia [f]
Jean-François Plante [a]
Juliana Schulz [a]
Marie-Hélène Simard [g]
Ernest Tafolong [a]
Véronique Tremblay [g]

[a] *HEC Montréal*
[b] *Université Concordia*
[c] *MAP5, Université Paris Cité*
[d] *ETS – École de Technologie Supérieure*
[e] *University of Regina*
[f] *Université de Montréal*
[g] *Beneva*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Aflaki, Bizimana, Cwiling, Farhangian, McVittie, Monsia, Plante, Schulz, Simard, Tafolong, Tremblay

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

2.1 Introduction

Beneva was created through the merger of La Capitale and SSQ Insurance, building on 75 years of experience. Beneva offers a range of insurance products, including auto, home, life, and commercial insurances, in addition to savings and investment products. Beneva has over \$26.8 billion in assets, 3.5 million clients and members, and over 5000 employees nationwide.

A fundamental task, which is a key element in several other analyses, is the modelling of customer lifetimes. This allows the company to develop an understanding of their clientele and acknowledge customer loyalty. In particular the lifetime of a customer is an important input in calculating a customer's lifetime value (CLV). The latter is determined according to the following equation:

$$CLV = \sum_{t=0}^{T} \frac{(R_t - C_t)S_t}{(1+i)^t} - AC, \tag{2.1}$$

where $R_t - C_t$ is the customer profitability at time t, S_t is the probability of customer retention at time t (i.e,. the survival probability at time t), $(1+i)^t$ is the discount rate, and AC is the customer acquisition cost.

The focus of this research work is to develop statistical techniques for modelling customer lifetimes, specifically, estimating the survival function $S(t)$. While there exists an extensive literature in the field of survival analysis, the peculiar features of the actual data and the business setting are such that traditional methods are not entirely adequate. Before exploring the statistical methods in detail, we will give a detailed description of the data and business constraints in Section 2. Section 3 will discuss some approaches considered, including parametric, non-parametric, and semi-parametric modelling, survival trees and forests, and a Bayesian approach. Concluding remarks are given in Section 4.

2.2 Data

A customer's lifetime is defined as the time spent as an active customer with the company. To determine the lifetime of a client, the starting event date is defined as the time at which the client purchases a policy with Beneva, and the failure event date is the time at which the client cancels his or her insurance policy. The lifetime of a client is defined as the difference between these two dates. Note that a customer who leaves and later returns to the company is treated as two separate (and independent) customers.

The data considered for the analysis consists of active and new customers between the years 2010 and 2021. More specifically the data includes all customers who had an active policy in 2010, as well as all new clients who purchased a policy between 2010 and 2021. This particular sampling procedure led to certain challenges that needed to be properly addressed in order to carry out a valid analysis. In particular the sample consists of two distinct "cohorts" of clients: one from the cross-section of clients who were active in 2010 and the other resulting from clients entering between 2010 and 2021.

Customers coming from the cross-section in 2010 were only partially observed: their starting time was accurately recorded, but only covariate information as of 2010 was available. Further note that because these individuals were cross-sectioned in 2010, the lifetimes of observed customers are not fully representative of the true lifetime distribution. This is due to the fact that individuals who purchased insurance policies prior to 2010 and then left prior to 2010 are not captured in the data at all. As a result, the lifetimes for the cross-sectioned customers are in fact left-truncated. In addition to the presence of left-truncation, we note that the failure events may not be fully observed for all individuals, as a particular failure event may occur past the 2021 cutoff. As such a portion of the lifetime data are also right-censored.

In order to model customer lifetimes adequately based on this data, we must account for the above issues in our statistical analyses. As such we defined the problem in terms of two distinct cohorts: a *prevalent* cohort and an *incident* cohort, which will now be described in more detail.

2.2.1 Insurance data in context of survival analysis

Define the set of individuals who purchased policies after 2010 and before 2021 as the *incident cohort*. For each individual, we observe not only their full covariate information (at baseline), but we observe their purchase date and their lifetime with the company up to 2021. The lifetimes of these individuals are commonly referred to as *right-censored failure time data*. We present a representation of these data in Figure 2.1.

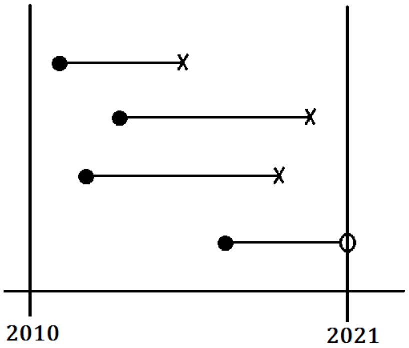

Figure 2.1: A graphical representation of the Beneva client lifetimes in the incident cohort. The solid circles represent the policy purchase events, the crosses represent the failure events and the open circles represent the right-censoring event

We refer to individuals who were cross-sectioned in 2010 as the *prevalent cohort*. For each individual that was sampled by the cross-section, we observe their policy purchasing date and their failure/censoring event dates. Covariate information is also observed for the prevalent cohort customers, however, these are measured at 2010 rather than at the policy purchase date (baseline). As discussed above, the lifetimes of the observed subjects captured by the cross-section are biased and so they are commonly referred to as *left-truncated right-censored failure times*. In Figure 2.2, we provide a representation of the failure/censoring times captured by the cross-section in 2010.

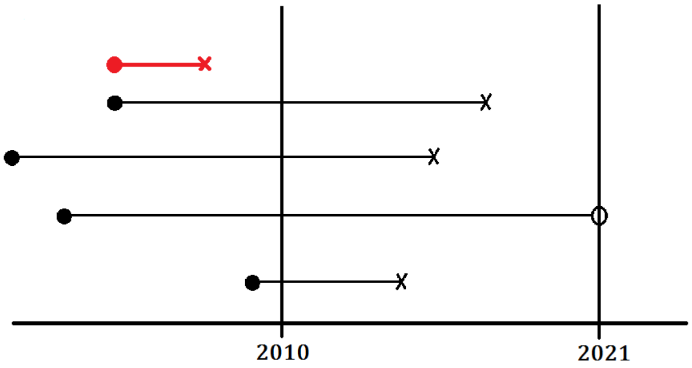

Figure 2.2: A graphical representation of the Beneva client lifetimes in the prevalent cohort. The solid circles represent the policy purchase events, the crosses represent the failure events and the open circles represent the right-censoring events. The subject with the red failure time would be unobserved as their purchase date occurred prior to 2010 and they left the company before 2010.

Each set of incident and prevalent cohort data has its own advantages and disadvantages with respect to the other. The incident cohort failure/censoring time data is sampled directly from the population under study and so no specialized techniques are required to adjust for bias, as with the data collected from the prevalent cohort. Conversely, since the incident failure/censoring times have their purchasing dates occurring between 2010 and 2021, the longest observed failure/censoring time in the cohort is bounded from above by 11. If the follow-up period were even shorter, the incident cohort data would provide little information on customers who stay with the company for long periods of time. The prevalent cohort failure/censoring times, however, are not restricted in length by the follow-up period since the purchase dates can occur arbitrarily far before 2010. Conversely, the observed failure/censoring times of the prevalent cohort will tend to be longer and so may not accurately model the survival function for customers who only stay with the company for a short period of time.

2.2.2 Missing values and time-varying covariates

In the data provided by Beneva, there are various forms of missing data. Firstly, and as most commonly encountered in practice, there are missing values that occur sporadically. In some instances the occurrence of missing values is in fact meaningful: in these cases an informed imputation was used (e.g., an "NA" in fact meaning 0). In the other cases standard imputation methods were used, assuming the values were missing at random.

An important limitation of the data available is that the covariate information is incomplete. In the data available for the workshop, covariate information is given at only one point in time for all individuals. For subjects in the incident cohort, covariates are those measured at the time of policy purchase (i.e., baseline). For the prevalent cohort, covariates are only available at the cross-sectioned time, i.e. 2010. In this sense the covariate information represents baseline information for the incident group, and rather lagged covariates for the prevalent cohort. Note that in reality, the majority of covariates used in actuarial models (including customer lifetime models) are in fact time-varying. This means that a trajectory of covariates would technically be available internally at Beneva for all clients in the incident cohort. For the prevalent cohort, however, limited data is available: while the origin date is known, covariate information prior to 2010 is unavailable. Thus the trajectories of longitudinal covariates are only partially observed (from 2010 onward) for the prevalent cohort. Certainly, a more realistic treatment of missing values would need to reflect the time-varying nature of the covariates. Due to the limitations imposed by the available data, we were unable to address this problem. Moreover, because longitudinal covariate trajectories are not available for any clients, a proper handling of the lagged covariates for the prevalent group was not feasible in the context of this workshop.

2.3 Approaches

There are a variety of different survival analysis approaches that can be used for modelling customer lifetimes within the context presented by Beneva. Note that the incident and prevalent cohort data sets could be modelled separately or combined. We chose the latter approach as it includes all available data in our analysis and the advantages of a specific cohort data set can make up for the weaknesses of the other cohort.

2.3.1 Statistical cohort approaches

To formalize the discussions above regarding incident cohort and prevalent cohort data sets, we introduce some mathematical notation. Let T_i, for $i = 1, 2, ..., n_{inc}$, denote the periods of time from purchasing the policy to leaving the company (i.e., the failure times). Some individuals may still have an active policy as of 2022, meaning their policy time T_i is *right-censored*. Let $C_i = 2022 - O_i$ for $i = 1, 2, ..., n_{inc}$ denote the right-censoring times for each of the n_{inc} subjects. When the policy purchase dates O_i are assumed to be independent of the holding times (i.e., the failure times), T_i, we may re-center all the holding/censoring times back to a common origin without affecting our inference. Thus, the observed incident cohort data set consists of the pairs

$$\mathcal{O}^{inc} = \{(X_i, \delta_i) = (\min(T_i, C_i), 1\{T_i < C_i\}) : i = 1, 2, ..., n_{inc}\}.$$

Using the same notation as for the incident cohort, let $O_i < 2010$ for $i = 1, 2, ..., k$ denote the purchasing date of the insurance policy for each subject occurring prior to 2010 and associate with each purchase date a holding time T_i for $i = 1, 2, ..., k$. If a cross-section of active policies is taken in the year 2010, the data will consist of a series of purchase dates O_j and holding times T_j for $j = 1, 2, ..., n_{prev}$, with $O_j + T_j > 2010$. Note that $n_{prev} <= k$ holds (i.e., the number of policy holders captured by the cross-section in 2010 is possibly less than the number of policy holders with policy purchase dates occurring before 2010) and that each individual j that was sampled has a policy holding time that is "long-enough" to be captured by the cross-section. As in the incident cohort setting, the holding times are right-censored at the end point of 2022. We typically refer to the periods of time $A_j = 2010 - O_j$ for $j = 1, 2, ..., n_{prev}$ as the *left-truncation times*. We note that each individual has their own left-truncation time as their policy purchase date is in fact random. As such our analyses must account for the distribution of the random truncation time A_j for the prevalent cohort. Thus the observed prevalent cohort data set consists of the triples

$$\mathcal{O}^{prev} = \{(A_j, X_j, \delta_j) = (A_j, \min(T_j, C_j), 1\{T_j < C_j\}) : T_j > A_j, j = 1, 2, ..., n_{prev}\}.$$

Throughout our analyses, we make the following assumptions regarding the incident and prevalent cohort data sets.

Incident Cohort Data Assumptions

1. The failure times $T_1, ..., T_{n_{inc}}$ are identically distributed according to some distribution function F.
2. The censoring times $C_1, ..., C_{n_{inc}}$ are identically distributed according to some distribution function H. No information regarding F may be gained from H (i.e. the non-informative censoring property) and each of the T_i and C_i are independent (i.e., the random censoring assumption).
3. The purchase times $O_1, ..., O_{n_{inc}}$ are independent of the failure times (i.e., the no cohort effect assumption).

Prevalent Cohort Data Assumptions

1. The underlying failure times $T_1, ..., T_{n_{prev}}$ are identically distributed according to some distribution function F.
2. The censoring times $C_1, ..., C_{n_{prev}}$ are identically distributed according to some distribution function H. No information regarding F may be gained from H.
3. The underlying left-truncation times $A_1, ..., A_{n_{inc}}$ are independent and identically distributed according to some distribution function G. No information regarding F may be gained from G.

Parametric marginal modelling

One approach to modelling the distribution for the policy holder duration period is to assume the failure time distribution is defined up to some unknown set of parameters and then estimate these parameters through a technique such as maximum likelihood estimation. We investigate this approach below and examine the estimates when applied to the Beneva dataset.

Suppose the underlying failure times of both the incident and prevalent cohort data sets have continuous distribution function $F(\cdot; \boldsymbol{\theta}) = 1 - S(\cdot; \boldsymbol{\theta})$ and density function $f(\cdot; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is some set of unknown parameters. The likelihood function for the incident cohort data set is given by:

$$\mathcal{L}_{inc}(\boldsymbol{\theta}) \propto \prod_{i=1}^{n_{inc}} f(x_i; \boldsymbol{\theta})^{\delta_i} S(x_i; \boldsymbol{\theta})^{1-\delta_i} \tag{2.2}$$

where the proportionality statement has removed the terms involving H that do not contain any information on $\boldsymbol{\theta}$ [5]. The corresponding incident cohort maximum likelihood estimator is given by $\hat{\boldsymbol{\theta}}^{inc}$, which maximizes the likelihood in Equation (2.2).

For a prevalent cohort data set, we must account for the feature that the observed failure/censoring times are sampled with a bias from the cohort under study. To account for this we include the term $\mathbb{P}(T > A; \boldsymbol{\theta})$ in the denominator of the likelihood function. Thus the likelihood function for the prevalent cohort data set is given by

$$\mathcal{L}_{prev}(\boldsymbol{\theta}) \propto \prod_{j=1}^{n_{prev}} \frac{f(x_i; \boldsymbol{\theta})^{\delta_j} S(x_j; \boldsymbol{\theta})^{1-\delta_j}}{\mathbb{P}(T > A; \boldsymbol{\theta})} \tag{2.3}$$

where the corresponding parametric maximum likelihood estimator is given by $\hat{\boldsymbol{\theta}}^{prev}$. We note that since policy holders may start their policies at any random time, the denominator of Equation (2.3) is not expressed as $S(a; \boldsymbol{\theta})$ for some fixed constant a. Given the inclusion of this complicated expression, we propose below a simplification based on the observed policy start dates from the incident cohort data.

Assuming the observed purchasing dates in the incident cohort have the same distribution as the purchasing dates occurring before 2010, we can use the information from the incident cohort purchase dates to infer the distribution of A. Observing the histogram of the purchasing dates in the incident cohort in Figure 2.3, we find that the distribution is approximately a discrete Uniform distribution. Thus we assume the policy purchase dates follow a discrete Uniform occurring as early as 1960. Using this approximation, we can simplify the denominator of likelihood (2.3) to:

$$\mathbb{P}(T > A; \boldsymbol{\theta}) = \sum_{i=1}^{50} \mathbb{P}(T > A; \boldsymbol{\theta} | A = i) \mathbb{P}(A = i) =$$

$$\frac{1}{50} \sum_{i=1}^{50} \mathbb{P}(T > i; \boldsymbol{\theta} | A = i) = \frac{1}{50} \sum_{i=1}^{50} \mathbb{P}(T > i; \boldsymbol{\theta}) = \frac{1}{50} \sum_{i=1}^{50} S(i; \boldsymbol{\theta}).$$

When the incident and prevalent cohort data sets are assumed to be independent of each other, we can combine the data to find the combined cohort parametric maximum likelihood estimator. The combined cohort likelihood function is given by:

$$\mathcal{L}_{comb}(\boldsymbol{\theta}) = \mathcal{L}_{inc}(\boldsymbol{\theta}) \times \mathcal{L}_{prev}(\boldsymbol{\theta}) \tag{2.4}$$

with corresponding parametric maximum likelihood estimator given by $\hat{\boldsymbol{\theta}}^{comb}$ [10].

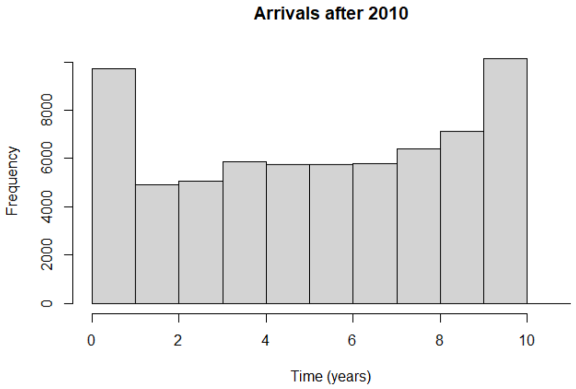

Figure 2.3: A histogram of the observed policy start dates measured from 2010 for the Beneva incident cohort.

Using the approaches discussed above, we fit multiple parametric models: Exponential, Weibull, and Gamma to the incident, prevalent, and combined cohort data sets. For brevity, we only present the survival function estimates for the Exponential model below in Figure 2.4.

From the fitted survival curve estimates, we found that the median policy holding time was approximately 7 years whereas the probability of having a policy of at least 5 years (resp. 10 years) was approximately 0.6 (resp. 0.35).

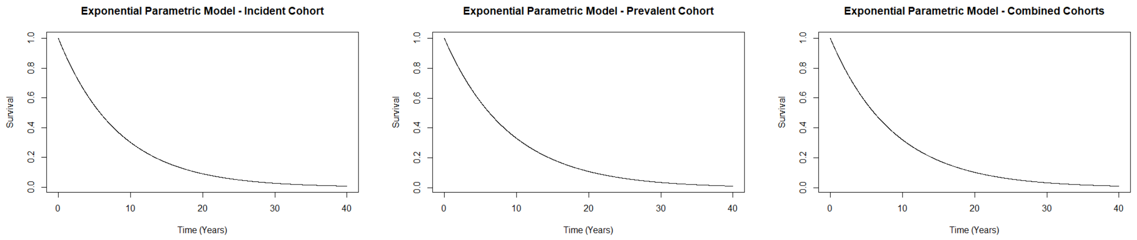

(a) **(b)** **(c)**

Figure 2.4: Exponential survival function estimates using the policy holding times for the incident (a), prevalent (b), and combined cohort (c) subjects.

Nonparametric estimation

One difficulty with using parametric models is that the model may be misspecified by the user yielding possibly flawed inferences regarding the underlying population. An alternative technique is to estimate the survival function non-parametrically through estimators that do not depend on the underlying failure time distribution.

Given a set of right-censored failure time data, the non-parametric maximum likelihood estimator of the survival function $S = 1 - F$ is given by the Kaplan-Meier estimator:

$$\hat{S}(t) = \prod_{i:t_i \leq t} \left(1 - \frac{d_i}{m_i}\right) \tag{2.5}$$

where t_i is an observed failure time, d_i is the number of failures that occurred at time t_i, and m_i is the number of individuals who have not yet failed or been censored up to time t_i. It can be shown that the estimator $\hat{S}$ is uniformly consistent and weakly convergent to a tight Gaussian process [1].

Given a set of left-truncated right-censored failure time data, the non-parametric estimator of S has a form similar to that of the Kaplan-Meier estimator. The non-parametric maximum likelihood estimator is given by:

$$\hat{S}(t) = \prod_{i:t_i \leq t} \left(1 - \frac{d_i}{m_i}\right) \tag{2.6}$$

where t_i is an observed failure time, d_i is the number of failures that occurred at time t_i, and m_i is the number of individuals who are at risk at time t_i. That is, m_i includes all those subjects with failure/censoring times, x_j, that are larger than or equal to the failure time t_i with truncation times a_j that are less than or equal to the failure time t_i (i.e., $m_i = \sum_j 1(a_j \leq t_i \leq x_j)$) [11].

When incident and prevalent cohort data are combined, the product-limit estimator of S can be easily obtained. By viewing the incident cohort data as prevalent cohort data with left-truncation times equal to 0, they may be automatically incorporated into the estimator given in (2.6) [8, 12]. The corresponding combined cohort product-limit estimator can also be shown to be uniformly consistent and weakly convergent to a tight Gaussian process.

The estimator in (2.6) was also fit to the Beneva data and the results are shown in Figure 2.5. For the incident, prevalent, and combined cohort survival curve estimates, we found the median policy holding times to be approximately six years, seven years, and six years, respectively.

There are both advantages and disadvantages to these non-parametric estimators. First, the survival function estimator for the incident cohort data set is restricted by the right-end point (2021) of the data collection window, meaning that the largest observed failure/censoring time is at most 11 years even though the survival function may be defined with non-zero probability past the period of 11 years. This restriction is not present in the prevalent cohort data as the policy purchase dates may occur at

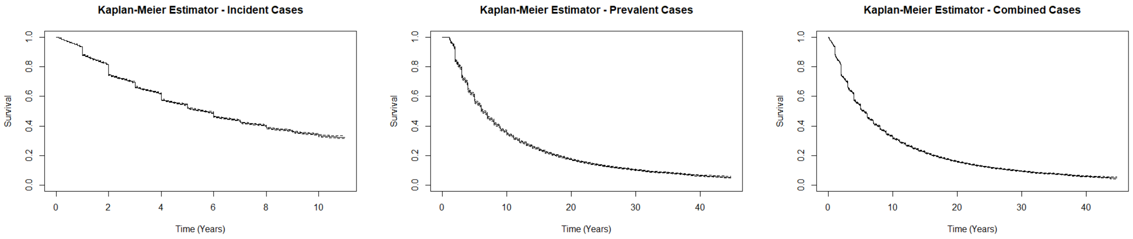

(a) (b) (c)

Figure 2.5: Kaplan-Meier survival function estimates (solid lines) and corresponding 95% confidence intervals (dashed lines) using the policy holding times for the incident (a), prevalent (b), and combined cohort (c) subjects.

any time prior to 2010, implying that the tail of the survival function is well estimated. When the data are combined, the survival function estimator incorporates both the short policy holding times from the incident cohort subset as well as the longer policy holding times from the prevalent cohort subset.

We also note that since the nonparametric maximum likelihood estimators are based directly on the observed data, the corresponding estimates have visible drops occurring every calendar year. We attribute these drops to the fact that the majority of customers will cancel upon renewal of their insurance policies at the end of their policy term. Generally speaking, the survival curves appear to suggest that policy holders are not likely to break and prorate their policy packages while their policies are still active and are more likely to wait until the policy term expires before leaving Beneva or purchasing a new policy.

Semi-parametric modelling

A quantity of interest for insurance providers is the risk that an individual will end their insurance policy with the company. The hazard function is the instantaneous rate of change in the conditional probability of failure at time t, as given by the formula below:

$$\lambda(t) = \lim_{dt \to 0} \frac{\mathbb{P}(T \in [t, t+dt) | T > t)}{dt}. \tag{2.7}$$

A modelling approach discussed by Cox relates the conditional hazard function (conditional on a set of covariates) to the product of a baseline hazard function (free of covariates) and an exponential risk function [2]. We define the hazard function of the Cox proportional hazards model below:

$$\lambda(t|\boldsymbol{Z}) = \lambda_0(t) \exp(\boldsymbol{Z}\beta). \tag{2.8}$$

For right-censored failure time data, the standard approach for estimating the β parameters is to maximize the partial likelihood function [1] given below:

$$\mathcal{L}^{inc}(\beta) = \prod_{i=1}^{n_{inc}} \left(\frac{e^{\boldsymbol{Z}_i \beta}}{\sum_{j: X_i \leq X_j} e^{\boldsymbol{Z}_j \beta}} \right)^{\delta_i}. \tag{2.9}$$

Similarly, for a set of left-truncated right-censored failure time data, the risk function β parameters can be estimated through a partial likelihood function accounting for the presence of left-truncation:

$$\mathcal{L}^{prev}(\beta) = \prod_{i=1}^{n_{prev}} \left(\frac{e^{\boldsymbol{Z}_i \beta}}{\sum_{j: A_j \leq X_i \leq X_j} e^{\boldsymbol{Z}_j \beta}} \right)^{\delta_i}. \tag{2.10}$$

Interestingly, there are some cases when there are ties in the observed failure/censoring times that the product of the partial likelihoods $\mathcal{L}^{inc}(\beta)\mathcal{L}^{prev}(\beta)$ is not fully efficient [9]. To counteract this possibility, we follow the same approach as in the nonparametric estimation setting by regarding the incident cohort data as prevalent cohort data with truncation times equal to 0. We can then incorporate all the incident cases into a single partial likelihood function with the prevalent cohort cases.

Due to the data privacy restrictions of Beneva, the covariates that were provided in the observed data set were not labelled. As a result of the anonymity of the covariates, parameter interpretations were not possible. Nonetheless we applied the proportional hazards model incorporating the available covariates. We also considered treating the observed truncation time as an additional covariate in the model. Specifically, using the proportional hazards model, we explored how the survival curve estimates depend on whether or not the observed truncation time is included as a covariate in the risk function. In Figure 2.6, we plot the estimated baseline survival functions when the covariate of the left-truncation time is included (or not included) as a potential covariate. From the preliminary plot, we find that the inclusion of truncation as a covariate shortens the estimated holding time of individuals who purchased insurance policies. Using the techniques described above, Beneva can apply our methodologies to incorporate their specific covariates into the model in a more appropriate manner. This would allow a better understanding of the covariate effects on insurance policy holding times, and how these covariates affect the risk that an individual may cancel his or her policy.

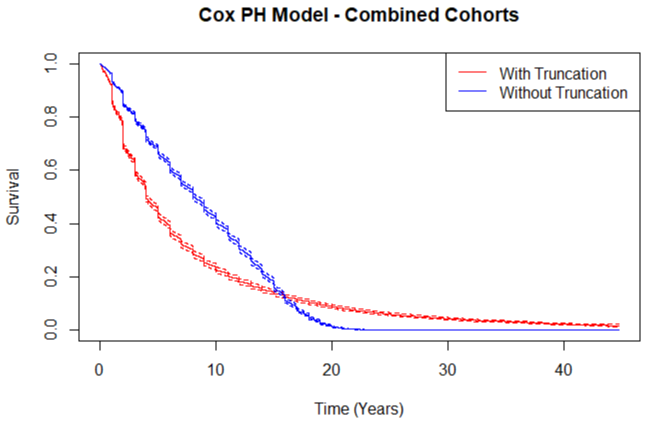

Figure 2.6: Survival curve estimates under the Proportional Hazards Model when the truncation time is/is not included as a covariate in the risk function.

2.3.2 Survival trees/forests

Ignoring the issue of incomplete time-varying covariate information, there are two particularities of the data that need special handling: first, the observed customer lifetimes are left-truncated right-censored, as previously described; secondly, there are sporadically missing covariate values, which we assume are missing at random. We will now explore the use of random trees and forests for modelling the customer lifetimes.

Random trees are a sub-category of binary trees where data are split into leaves according to a defined criterion. The process is repeated successively on every leaf until a stopping criterion is verified, such as a minimum number of data points in the leaf or when the splitting does not improve prediction. Growing a random forest consists of growing several random trees. Each one is trained on a random subset of the training data and the associated features. A cumulative hazard function can then be estimated on each tree. By averaging them, we obtain the ensemble cumulative hazard function [4].

The key to growing a random tree is the choice of splitting rule. In a more standard context without the presence of censoring or truncation, the Gini criterion and the entropy criterion are widely used. To account for right censoring, the logrank criterion [4] or a Poisson regression based on the log-likelihood [7] can be considered. The logrank criterion and the log-likelihood can be adjusted to account for left truncation. For more details about the latter, see Section 2 in [3].

In this application three approaches were considered using random survival trees and forests.

1. Construct a random survival forest that can handle left-truncated right-censored data and missing data simultaneously.
2. Create a new covariate accounting for left truncation. Then construct a random survival forest for right-censored data and missing data simultaneously, treating the left truncation as a feature.
3. Impute the missing data, and then fit the model to the completed data by growing a random survival forest for left-truncated right-censored data.

Method 1

There exists a random survival forests algorithm for right-censored data, which is capable of imputing missing data as the trees are grown [4]. This method, however, is not designed to take left truncation into account. An alternative approach for building random survival forests given in [13] can handle both left truncation and right censoring, but it does not impute missing values. It seems that there is not yet a readily available implementation combining both algorithms. Note that it may be feasible to incorporate the splitting rule from the latter algorithm into the former approach, although this requires further investigation.

Method 2

Another approach is to construct a new covariate accounting for left truncation. This covariate is set equal to the lifetime of the customer at the 2010 cross-section. More specifically, the truncation covariate is $[2010 - O_j]_+$, i.e., it is equal to 0 for clients who arrived after 2010 and $2010 - O_j$ otherwise. We can then grow a random survival forest for right-censored data that can also handle missing data. This induces a bias, however, as it does not properly capture clients who purchased insurance policies and left prior to 2010. Moreover the running time of the algorithm is quite large, especially with the amount of data considered.

Method 3

As an alternative approach, we considered imputing the missing data in an initial step and then using the completed data to fit the random survival forest for left-truncated right-censored data as given in [13]. While the algorithm of [13] is available in an R package, we ran into some implementation difficulties. To simplify, we considered an alternative package which consists of an extension of the CART algorithm adapted to left-truncated right-censored data [3], that is, a single tree is grown instead of a forest. The results of this algorithm can be a Kaplan-Meier curve for each terminal node as illustrated in Figure 2.7. The predicted survival curve of a new client is then given by the estimated Kaplan-Meier curve at the terminal node corresponding to their set of covariates.

2.3.3 Bayesian approach

As a final method, a Bayesian approach was explored. Following the work of [6], we focused on estimating the survival function assuming a Weibull distribution wherein the underlying likelihood function was appropriately adapted for left-truncated right-censored observations. Due to time constraints, a very preliminary analysis was carried out wherein the shape parameter was assumed known (held fixed at the maximum likelihood estimate obtained from a parametric approach). In this setting, the scale parameter has a conjugate gamma prior, thus leading to a simple closed-form expression for the Bayes estimator. Due to the time constraints of the workshop, our analysis ended here.

A more thorough investigation of Bayesian inference for modelling customer lifetimes is necessary. First one could consider distributions other than the Weibull model. In the case of the Weibull distribution (and other distributions as well), a more realistic analysis would consider estimating all

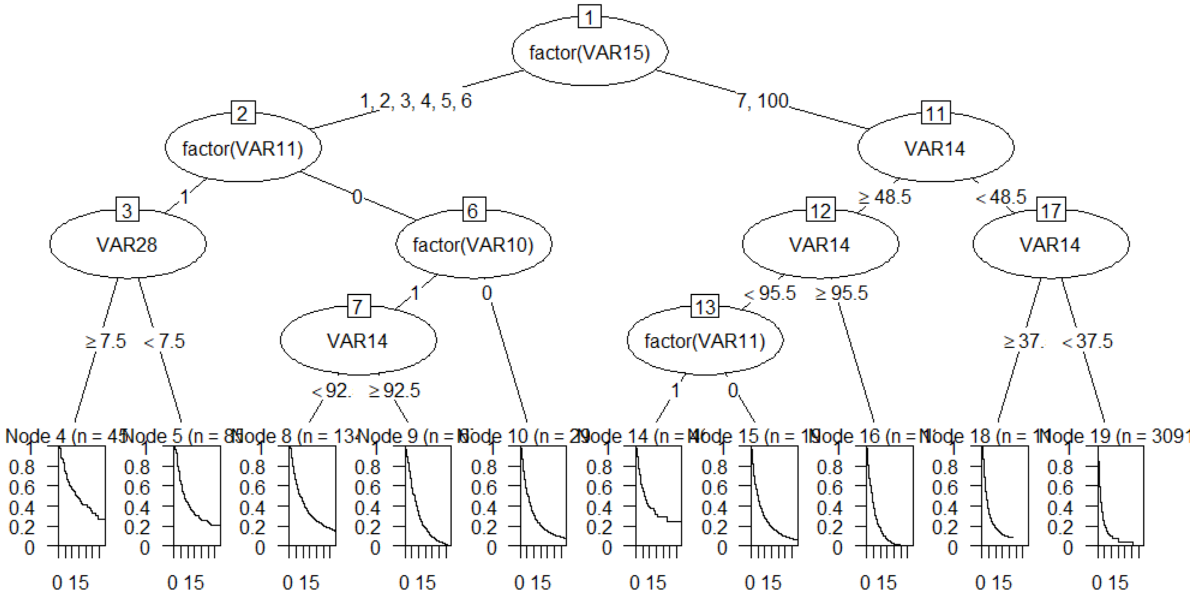

Figure 2.7: A random survival tree grown using the extension of the CART algorithm to left-truncated right-censored data. Missing data were imputed to yield completed data, which were then input into the algorithm. Each terminal node results in a Kaplan-Meier curve fit to data within the node.

model parameters, rather than assuming that some are fixed. In the context of customer lifetimes, the models should incorporate covariate information. As a first step, baseline (fixed) covariates could be considered while further explorations could consider including time-varying covariates. Finally, as detailed in Section 2.3.1 (see parametric marginal modelling), the underlying likelihood should be adapted to reflect the random truncation times for customers from the prevalent cohort.

2.4 Discussion

The approaches discussed in this report provide some preliminary analyses towards modelling left-truncated right-censored customer lifetimes arising from incident and prevalent cohorts. As demonstrated in the report, a proper handling of observations stemming from the two distinct cohorts allows for a combined analysis using all policies simultaneously. An overview of various estimation methods has been discussed, notably parametric, non-parametric and semi-parametric models, survival trees, and Bayesian inference. Of course all of the approaches considered in this report can be further explored with added layers of complexity: notably, a thorough investigation of the implementation of survival forests, alternative imputation methods for handling missing data, and a more rigorous Bayesian analysis.

It is important to note that while we have explored methods that appropriately handle the incident and prevalent cohorts, with left-truncated right-censored customer lifetimes, we were unable to address the issue of incomplete time-varying covariate information. Recall that in reality almost all covariates are indeed time-varying: in the observed data, however, covariates were only measured at one point in time. While baseline covariates are available at the time of policy inception for the incident cohort, only lagged covariate measurements are available for the prevalent cohort at the 2010 cross-section. Due to the constraints of the available data, we were simply unable to address this issue adequately. In a more detailed dataset, including time-varying covariates, imputation methods could be explored allowing one to impute missing covariate trajectories for subjects in the prevalent cohort; more sophisticated methods could also be used for imputing covariate values missing at random.

Acknowledgments

We would like to thank Beneva for submitting this problem and Odile Marcotte and Nancy Laramée for their organization and planning of the Twelfth Montreal Industrial Problem Solving Workshop.

Bibliography

[1] P. K. Andersen, Ø. Borgan, R. D. Gill, and N. Keiding. Statistical Models Based on Counting Processes. Springer Series in Statistics. Springer-Verlag, New York, NY, USA, 1993.

[2] D. R. Cox and D. Oakes. Analysis of Survival Data. Monographs on Statistics and Applied Probability. Chapman and Hall CRC, 1984.

[3] V. Fu and J. S. Simonoff. Survival trees for left-truncated and right-censored data, with application to time-varying covariate data. Biostatistics, 18(2):352–369, 2017.

[4] H. lshwaran, U. B. Kogalur, E. H. Blackstone, and M. S. Lauer. Random survival forests. The Annals of Applied Statistics, 2(3):841–860, 2008.

[5] J. D. Kalbfleisch and R. L. Prentice. The Statistical Analysis of Failure Time Data. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ, USA, 2 edition, 2002.

[6] D. Kundu and D. Mitra. Bayesian inference of Weibull distribution based on left truncated and right censored data. Computational Statistics and Data Analysis, 99:38–50, 2016.

[7] M. Leblanc and J. Crowley. Relative risk trees for censored survival data. Biometrics, 48:411–425, 1992.

[8] J. H. McVittie, D. B. Wolfson, and D. A. Stephens. A note on the applicability of the standard nonparametric maximum likelihood estimator for combined incident and prevalent cohort data. Stat, 9(1):e280, 2020.

[9] J. H. McVittie, D. B. Wolfson, and D. A. Stephens. A note on the partial likelihood estimator of the proportional hazards model for combined incident and prevalent cohort data. Accepted in Metrika, 2022.

[10] J. H. McVittie, D. B. Wolfson, D. A. Stephens, V. Addona, and D. Buckeridge. Parametric models for combined failure time data from an incident cohort study and a prevalent cohort study with follow-up. International Journal of Biostatistics, 17(2):283–293, 2020.

[11] M. C. Wang. Nonparametric estimation from cross-sectional survival data. Journal of the American Statistical Association, 86(413):130–143, 1991.

[12] D. B. Wolfson, A. F. Best, V. Addona, J. Wolfson, and S. M. Gadalla. Benefits of combining prevalent and incident cohorts: An application to myotonic dystrophy. Statistical Methods in Medical Research, 28(10-11):3333–3345, 2019.

[13] V. Yao, H. Frydman, D. Larocque, and J. S. Simonoff. Ensemble methods for survival function estimation with time-varying covariates. Statistical Methods in Medical Research, 20, 2022.

3 Prototype of a statistical dynamical hybrid model for high-resolution hydrodynamic simulations

Julie Carreau [a]
Anton Iatcenko [b]
Silvia Innocenti [c]
Pascal Matte [c]
Philippe Petitclerc [d]
Nassim Razaaly [e]

[a] *Polytechnique Montréal*
[b] *Simon Fraser University*
[c] *Environment and Climate Change Canada*
[d] *Université Laval*
[e] *Institut Pprime*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Carreau, Iatcenko, Innocenti, Matte, Petitclerc, Razaaly

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

3.1 Context and motivation

Environment and Climate Change Canada (ECCC) is developing operational environmental forecasting tools to characterize and monitor the evolution of different components of the Earth System, 24 hours a day, 7 days a week. These forecasting systems typically include a cascade of coupled models representing atmospheric, oceanic, ice, hydrological, hydrodynamic, and ecosystem conditions, from the global scale to the regional or even local scale. In lakes and rivers these systems are used as integrated tools for decision-making and water management, as aids to navigation, or in support of search and rescue operations and environmental emergencies. As such, research and development focuses on improving ECCC's prediction capabilities and capacity to deliver and sustain environmental prediction products and services to Canadians, researchers, industry, and government. Several applications require high-resolution (spatial and temporal) modelling of critical variables such as water depth or current velocity, particularly in coastal regions or near infrastructures.

Despite the availability of high-resolution data sets and high-performance computers, high-resolution simulations remain very expensive and the resulting computation times are sometimes inadequate for the application considered. Thus lower-cost solutions are generally computed by decreasing spatial and temporal resolutions, or reducing the forecast horizon, or simplifying the physical problem. A hybrid statistical-dynamical solution is sought in order to broaden the model applicability by improving the physical fidelity and increasing, at reduced cost, the computational speed and spatiotemporal resolution of the simulated hydrodynamic variables. Such a solution would also allow issuing longer-term forecasts (deterministic and ensemble-based) or climate projections. To accomplish this task, several emulators, or surrogate models, can be set up, using approaches borrowed from artificial intelligence and statistical or numerical methods aimed at reducing the dimensionality of the simulated problem. These techniques could, for example, allow us to generate high-resolution fields from low-resolution simulations, or to approximate fine-scale spatiotemporal variability using pre-simulated (simplified or dimension-reduced) solutions.

3.2 Problem description

As part of the Great Lakes Protection Initiative (GLPI), ECCC was given the task to model current and future water level fluctuations in the Great Lakes subject to various atmospheric, hydrological, and climatic conditions, with the aim of producing impact analyses of water level fluctuations (e.g. frequency and duration of flood cycles) on the distribution and evolution of coastal wetlands, including climate projections for a large number of future climate scenarios. This type of application requires a large number of simulations at a high spatial resolution, particularly near the coast, at a sufficiently small time step (sub-hourly), in order to capture high-frequency variations, and over long periods (30+ years). In order to reduce computational costs and the number of simulations to produce, various simplifications have been made. In particular two-dimensional (2D) simulations were performed, neglecting lake stratification and mixing; steady-state (time-invariant) simulations were conducted, using uniform and static scenarios to reconstruct temporal variability, neglecting high-frequency phenomena such as seiche oscillations; some variables were neglected, namely astronomical tides, atmospheric pressure, and ice, which are not modelled but identified as potentially having an impact on water level variations.

As part of the 12th Montreal Industrial Problem Solving Workshop (held in August 2022), the current work had the objectives of identifying existing approaches and exploring various strategies to improve the physical fidelity of the solutions, from forcing conditions to modelled water levels, and reducing the cost of long-term hydrodynamic simulations in the Great Lakes.

3.3 Case study: Lake Erie

Lake Erie is the shallowest of the Great Lakes with an average depth of 19 m (Figure 3.1). Because the influence of a strong, shallow stratification in Lake Erie remains limited, this lake can be approximated by a one-layer (2D) model. Although this assumption may be violated under certain circumstances, particularly in the deep eastern basin, wind-induced surface motions modelled at the basin scale remain mostly unaffected. Hydrodynamic simulations were therefore performed under time-varying conditions to account for the transient nature of this system. These real-time solutions thus offer a complete and realistic description of the wind-induced physical processes, i.e., wind set-ups and surface seiches. This method, however, is by far the most demanding in terms of computational and storage resources; for a one-year simulation, over 5000 lake-level maps of Lake Erie are produced and stored.

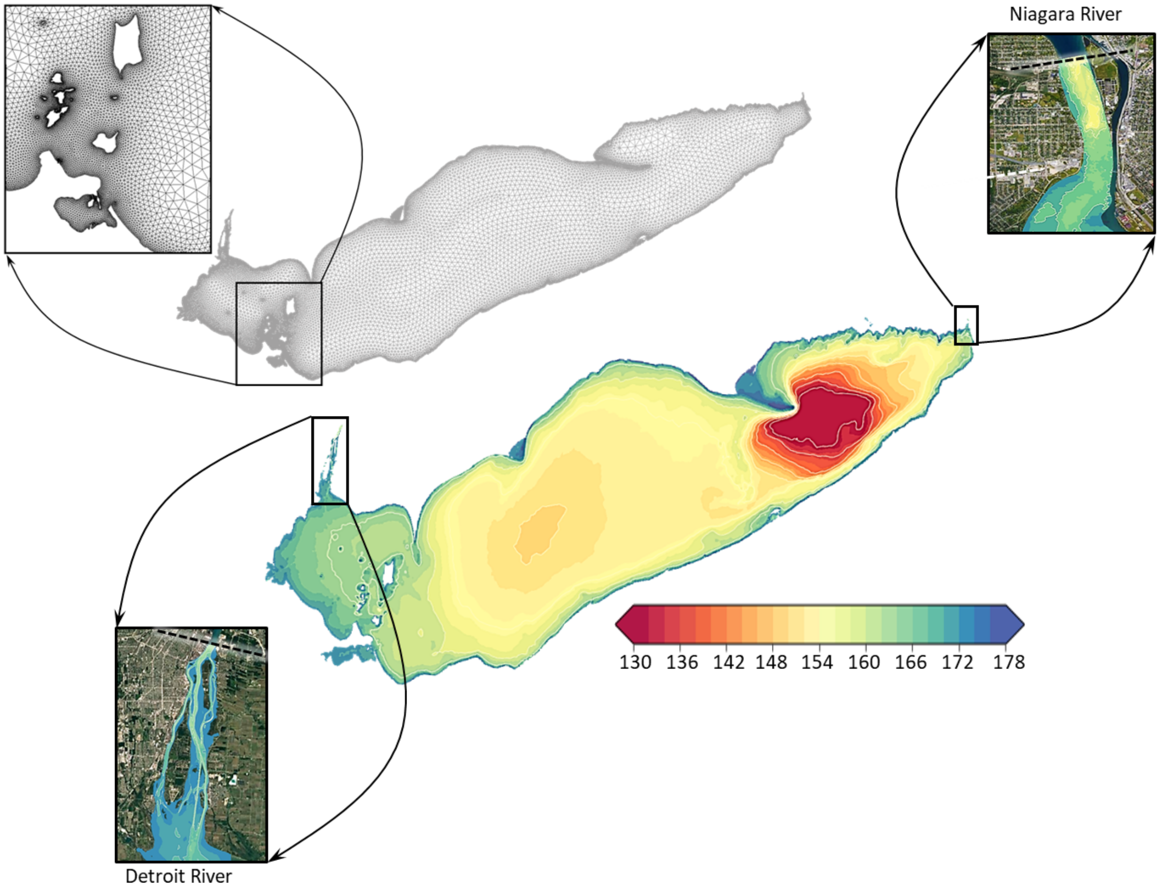

Figure 3.1: The mesh (top) and the bathymetry (bottom) of the Lake Erie hydrodynamic model. Insets show details of the lake's inlet (Detroit River) and outlet (Niagara River).

In contrast, for the other, much deeper, Great Lakes, a second modelling approach was developed, seeking to reproduce the physics of free-surface motions as observed in the mixed (upper) layer only. As the motion initiated by the atmospheric forcing is restricted to the surface Ekman layer whose depth varies with the wind intensity, the adopted strategy was to simulate only the surface layer dynamics by artificially modulating the lake bathymetry. In these cases the water column depth is used to estimate the wind effect within a 2D parametrization. As such, a calibration of the "effective" depth under different wind conditions was performed, considering both direction and speed. In this approach only static (i.e., steady-state) solutions can be obtained to avoid a violation of conservation of mass. In order to reproduce the transient water level fluctuations, the time series are reconstructed from pre-computed solutions rather than simulated continuously (i.e., using a time-stepping approach). Since it is not realistic to model all possible forcing combinations, a solution space was developed and used to estimate the state of the system, i.e., the tilt of the lake surface once in equilibrium with the wind. In this representation, only the initial deflection of the lake surface is modelled: the surface seiches initiated by the relaxation of the lake surface are not parametrized since a seiche is an oscillatory process resulting from an instability. Part of the energy will therefore be missing from the final time series reconstructed with the steady-state solutions.

During the temporal reconstruction, further simplifications are made on the gridded wind fields in order to match them with the modelled set of wind scenarios. In the scenario-based approach, the wind used to trigger a basin-scale motion is first obtained by low-pass filtering (with a 2-day cut-off period) the zonal, u, and meridional, v, wind components. Second, a time series was created from the gridded wind data sets by reducing the spatial distribution to a regional description of the wind, assumed to be representative of the predominant wind condition experienced (on average) by the basin. To do so a Gaussian Mixture Model was applied, using a fixed number of three clusters to describe lake-wide spatial variations. In Lake Erie, the availability of both a continuous (transient) solution and a scenario-based reconstruction allows the investigation of the impact of this simplification, given that only the static solutions are available in the other Great Lakes.

3.3.1 Proposed solutions

In the following sections, the three solutions proposed during the workshop will be presented. These solutions are to be seen as preliminary steps towards a global solution. Each solution addresses a distinct aspect of the problem, namely:

1. Seeking an analog method for the wind regionalization using Principal Component Analysis (PCA) projections instead of a single lake average, from which new static scenarios could be defined using information on wind patterns [Section 3.4];
2. Extracting modes of variability in water-level time series using Variational Mode Decomposition (VMD) and building a statistical model to predict high-frequency lake oscillations (e.g. seiches) from low-frequency events (e.g. wind setups) [Section 3.5];
3. Scaling and simplifying the hydrodynamic model (shallow-water) equations in order to obtain a low-cost model to be used for continuous, long-term simulations [Section 3.6].

3.3.2 Data

The following three data sets were available for preliminary tests during the workshop [github.com]:

- Hourly water level series from NOAA (Water Levels: Great Lakes Environmental Research Laboratory) and ECCC stations on Lake Erie:
  - series at six Canadian Hydrographic Service stations (Fisheries and Oceans Canada) and nine NOAA stations (NOS Center for Operational Oceanographic Products and Services);
  - mean lake level series estimated as the seasonally-adjusted average of the hourly water level series recorded at five stations [Port Stanley and Port Colborne, ON, Canada; and Toledo, Cleveland, and Fairport, OH, USA] following the standard NOAA methodology;
- Hourly wind series from the Climate Forecast System Reanalysis (CFSR): u et v [ms^{-1}] components of the wind at 10 m over a grid with 0.312° spatial resolution (approx. 38 km);
- Water levels simulated by the H2D2 hydrodynamic model [4, 5] and re-gridded to a 500m-resolution spatial grid with a temporal resolution of 24 min.

Two periods of two months were used for calibration: 15/3/1991–15/5/1991 and 15/3/1992–15/5/1992 were the validation periods.

3.4 Improvement of Scenarios: Principal Component Analysis (PCA) of winds

For this solution, the aim was to summarize wind fields by performing a projection onto a lower dimensional space. The wind fields are thus summarized as features consisting in their coordinates in that lower dimensional space. The low dimensional features should be related to the spatial patterns of the water level. Indeed the goal is to reconstruct the water level field based on the observed features of the wind fields. The reconstruction technique considered here is a simple analog method that boils down to seeking in the historical data time steps that are similar in terms of wind field features.

To perform the dimensionality reduction, we used Principal Component Analysis (PCA), a widespread statistical method that projects high dimensional data onto a lower dimensional hyperplane.

Let $(\boldsymbol{u}_t, \boldsymbol{v}_t) \in \mathbb{R}^D \times \mathbb{R}^D$ be the zonal and meridional wind components for a given hour t. PCA is applied separately on each wind component since the wind components are uncorrelated and their distribution is approximately Gaussian. Let $d \leq D$ be the dimension of the PCA hyperplane. Then PCA computes the following projections to obtain the features $\boldsymbol{\phi}_t$ and $\boldsymbol{\psi}_t$ in $\mathbb{R}^d$:

$$\boldsymbol{\phi}_t = A\boldsymbol{u}_t, \tag{3.1}$$

$$\boldsymbol{\psi}_t = B\boldsymbol{v}_t, \tag{3.2}$$

where A and B are $d \times D$ matrices whose rows are eigenvectors sorted in order of decreasing eigenvalue, resulting from the eigenvalue decomposition of the empirical covariance matrices of each wind component.

To examine the interpretability of the PCA projections, we define a training set consisting of the hourly wind components over Lake Erie at approximately 0.33° spatial resolution over the period 01/03/1991–30/04/1991. PCA is applied, as explained above, by computing the eigenvalue decomposition of the covariance matrices over the training set. The principal eigenvectors represent the main spatial patterns present in the training set and carry the spatial information. The eigenvectors associated with the two highest eigenvalues are presented in Figure 3.2, where the zonal and meridional components are combined to yield wind direction (indicated by arrows) and magnitude (indicated by the color scale). The first eigenvector (left panel of Figure 3.2) represents the spatial pattern that explains the largest fraction of variability in the data (about 90%). As the direction is constant and the wind magnitude varies only mildly, this spatial pattern is consistent with the original approach employed at ECCC in which wind fields are summarized with a single average value. The second eigenvector (right panel of Figure 3.2) represents a much more contrasted spatial pattern explaining the second largest fraction of variability in the data (about 5%). The temporal variability is contained in the feature vectors $\boldsymbol{\phi}_t$ and $\boldsymbol{\psi}_t$ (see (3.1)–(3.2)), as illustrated in Figure 3.3.

To evaluate the relevance of the features obtained by PCA for summarizing the wind fields, we designed the following procedure. The dimension of the PCA subspace is set to $d = 2$ for each wind component and the features of the zonal and meridional components are concatenated into a single feature vector $\boldsymbol{z}_t = (\boldsymbol{\phi}_t, \boldsymbol{\psi}_t) \in \mathbb{R}^4$. A test set is defined similarly as the training set but over the period 15/03/1992 - 30/04/1992 (i.e., a year later). The wind fields of the test set are projected onto the hyperplane of dimension two defined by the PCA eigenvectors computed over the training set. For each time step t' of the test set, let $\boldsymbol{z}_{t'} \in \mathbb{R}^4$ be the associated feature vector. To reconstruct the water level field, we look into the training set for the nearest neighbour in terms of feature vector. Namely let t^*, a time step in the training period, be chosen so that the following equation holds.

$$t^* = \arg\min_t ||\boldsymbol{z}_t - \boldsymbol{z}_{t'}||^2 \tag{3.3}$$

Then the water level field $\boldsymbol{h}_{t'}$ at time step t' is estimated by the water level field of its nearest neighbour as defined in the equation above from the training set.

$$\hat{\boldsymbol{h}}_{t'} = \boldsymbol{h}_{t^*} \tag{3.4}$$

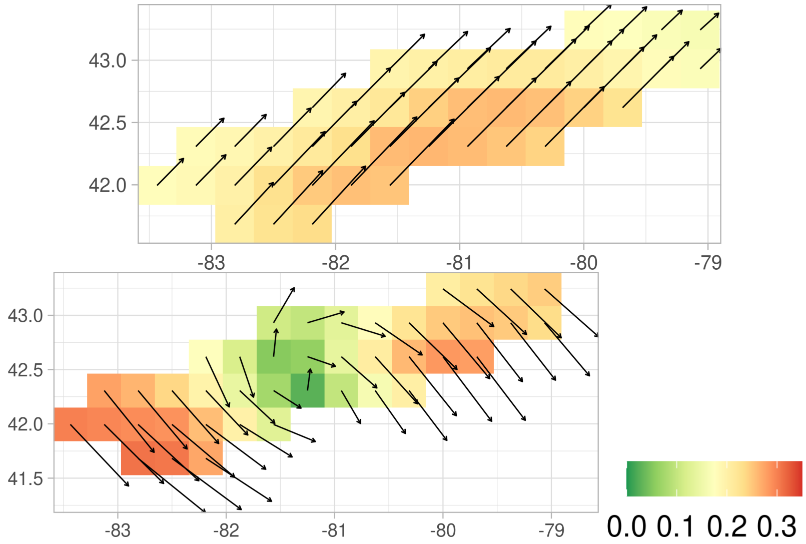

Figure 3.2: Wind direction (arrows) and magnitude (color scale) of the two principal eigenvectors (from left to right) of the zonal and meridional components obtained from PCA applied on the hourly wind fields over the period 01/03/1991–30/04/1991.

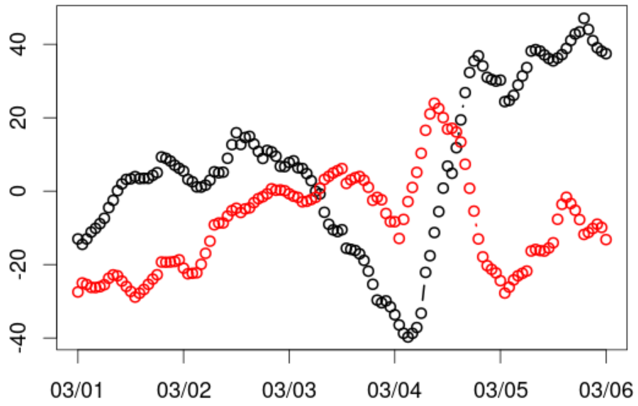

Figure 3.3: Temporal evolution of the first element of each feature vector, ϕ_t (in black) and ψ_t (in red), over the first five days of the training set (01/03/1991–30/04/1991).

The performance of the nearest neighbour reconstruction is measured by the Root-Mean-Square Error (RMSE), defined as

$$\sqrt{\frac{1}{T'}\sum_{t'=1}^{T'}\left(\boldsymbol{h}_{t'}-\hat{\boldsymbol{h}}_{t'}\right)^2} \tag{3.5}$$

where T' denotes the number of time steps in the test set.

The RMSE yielded by the approach based on the first two features of PCA applied to each wind component over the test set is compared to the RMSE yielded by the original ECCC approach in Figure 3.4. As can be seen from the RMSE, the PCA approach provides a performance similar to that of the original ECCC approach. There are likely at least two reasons to explain this result: (1) the reconstruction approach is not the one employed by ECCC (we could not use it as it involved running the hydro-dynamical model) and (2) PCA features could be improved by resorting to nonlinear dimensionality reduction techniques and/or to techniques accounting for extreme events.

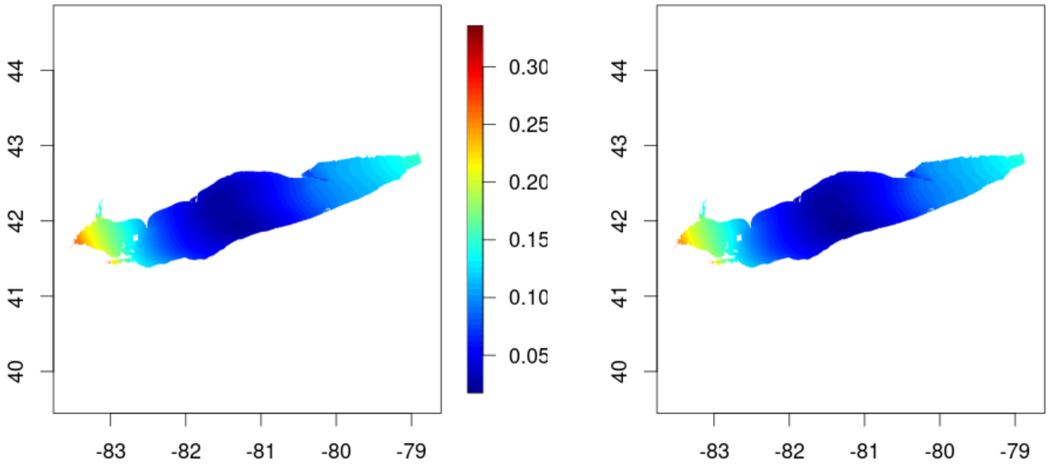

Figure 3.4: RMSE over the test set for the original ECCC approach (left panel) and for the $d = 2$ PCA features (right panel).

3.5 Seiche scenario construction: Variational Mode Decomposition (VMD) of the water levels

Lake water levels respond to a series of atmospheric and hydrological influences such as wind and river flow. The objective of this section is to analyze observed water levels at contrasting gauges around Lake Erie in order to understand their different modes of variability and to identify potential relationships between the low-frequency modes (e.g. associated with wind setup) and the high-frequency modes (e.g. associated with seiches). If such a relationship exists, the idea would be to build a predictive model capable of reconstructing the seiche signals in a synthesized way (e.g. amplitudes, phases) from identified predictors such as setup or wind. These predictors could come directly from the low-frequency modes of observed or simulated water levels or, alternatively, from wind fields or their principal components obtained by PCA.

Variational Mode Decomposition (VMD) [1] was applied to separate the water level signals into different modes of variability. Two contrasting stations in Lake Erie were chosen, namely Toledo (Ohio, USA) and Port Stanley (Ontario, Canada) (Figure 3.5). Water levels at Toledo, located at the western end of Lake Erie, are the most impacted by wind setup/set-down events, as illustrated in Figure 3.5. While the water level anomaly reached up to 1.5 m in 1991 at Toledo, it only reached 0.4 m at Port Stanley, located on the North Shore midway from the eastern and western ends of the lake. This reveals how wind events predominantly oriented along the lake principal axis create a slanted lake surface affecting local water levels, most importantly at the lake extremities. These effects are even more pronounced in the western portion of the lake due to shallower bathymetry [3]. For the remaining of this section, results will be presented only for the Toledo station, which presents the largest variability in water levels.

A total number of 26 modes were extracted using VMD, allowing for a separation of the water level signals into the dominant modes (or frequencies). A sample (the first 16 modes) is provided in Figure 3.6 for the Toledo station. In a similar fashion the corresponding amplitude spectrum obtained from the Fast Fourier Transform (FFT) is presented in Figure 3.7. It can be seen that the VMD method effectively extracts and decomposes the signal variability into distinct modes with minimal overlap between neighbouring frequency bands. In addition, the method completely avoids mode mixing, meaning that the energy contained in each mode is limited to a narrow band. This represents an advantage over other techniques, such as empirical mode decomposition, and is a required feature if a predictive model is to be developed.

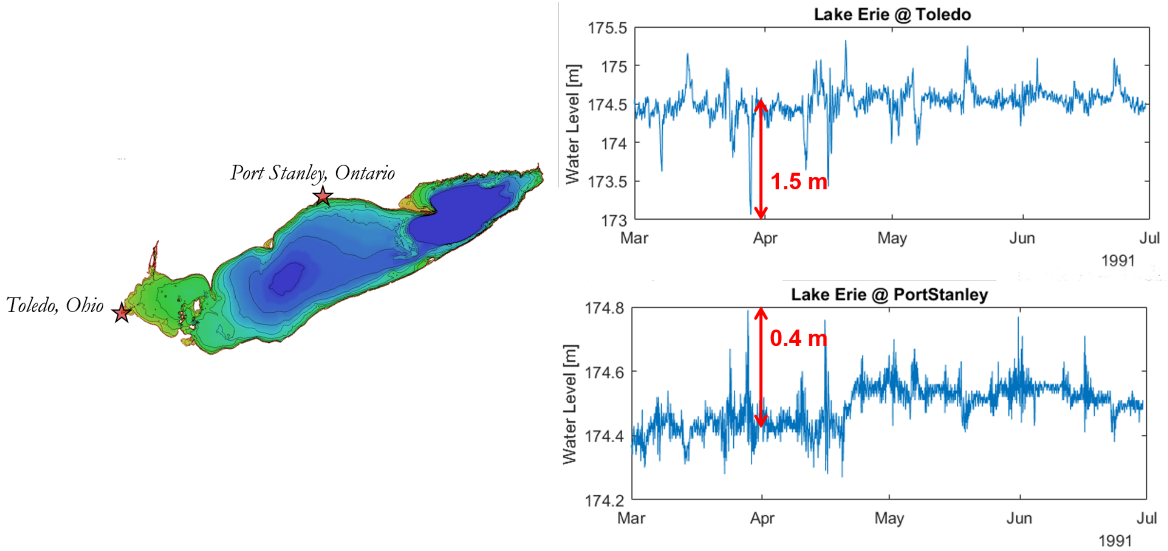

Figure 3.5: (left) Bathymetry of Lake Erie; (right) observed water level time series in Lake Erie at Toledo (Ohio) and Port Stanley (Ontario).

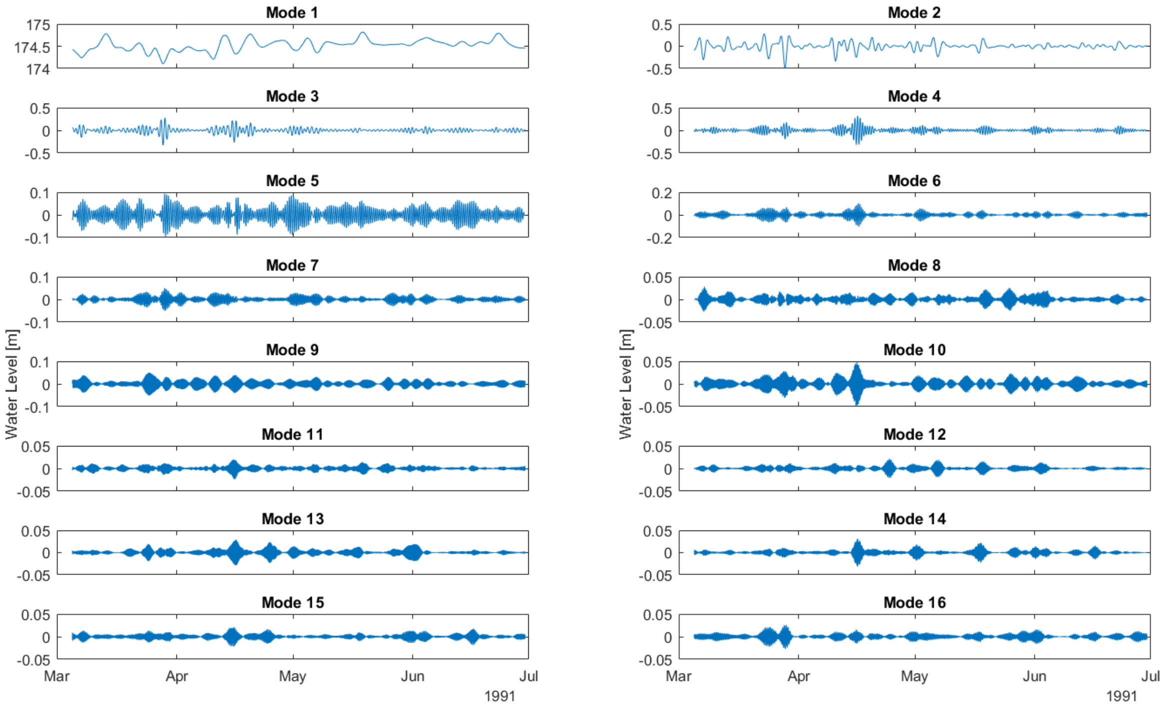

Figure 3.6: The first 16 decomposed modes from VMD at the station of Toledo (Ohio) in Lake Erie.

The FFT amplitudes for each mode are compared with the original spectrum at Toledo in Figure 3.8. Tides and expected seiche frequencies [7] are also identified. Relevant low and higher frequencies are well separated by VMD and can be represented by one or a few modes. As such, Figure 3.9 presents time series reconstructions for selected groups of frequencies, from the low-frequency band associated with wind setup events to the diurnal and higher frequency bands associated with a mix of tidal influence and seiche oscillations. By visual inspection of Figure 3.9, it can be seen that setup events are correlated to some degree with higher frequency oscillations. An attempt has thus been made to

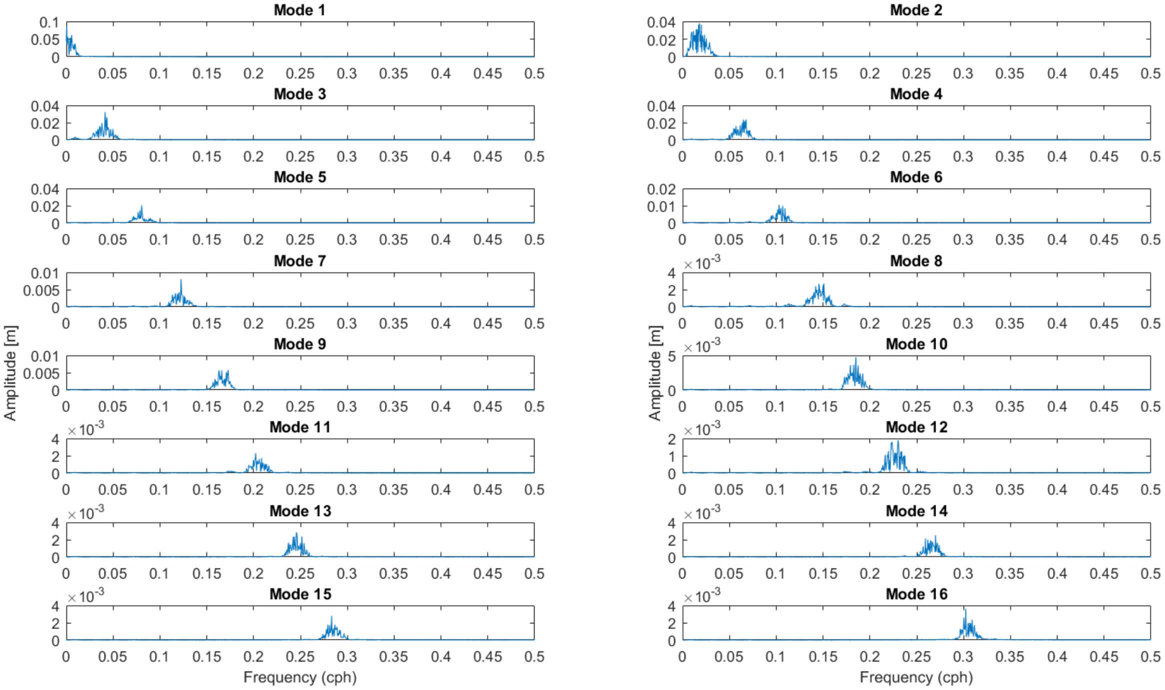

Figure 3.7: Amplitude spectrum of the first 16 VMD modes at the station of Toledo (Ohio) in Lake Erie.

construct a simple regression model to predict high-frequency oscillations (h_{HF}) from the low-frequency signal (h_{LF}). The model can be expressed as follows:

$$h_{HF}(t) = \beta_0 + \sum_{j=1}^{N_{\text{modes}}} (\beta_{1,j}(t)\cos\omega_j(t) + \beta_{2,j}(t)\sin\omega_j(t)) + \epsilon, \tag{3.6}$$
$$\text{where } \beta_{i,j}(t) = a_{0,i,j} + a_{1,i,j} \times h_{LF}(t),$$

and t denotes the time, $a_{n,i,j}$ the model parameters, ω_j the mode centre frequencies, and ϵ the error term. For each mode the signal amplitude and phase can be computed as follows.

$$A_j(t) = \sqrt{(\beta_{1,j}^2 + \beta_{2,j}^2)} \qquad \text{and} \qquad \phi_j(t) = \tan^{-1}(\beta_{2,j}/\beta_{1,j}) \tag{3.7}$$

This model was applied using a low-frequency signal, h_{LF}, defined as a 48-h moving standard deviation of the 2-day setup signal (mode 2). Water level predictions for the diurnal mode and 14-h seiche mode are shown in Figure 3.10. Predictions are better at the diurnal frequencies than at higher frequencies. It is not sufficient, however, to use a representation of the setup signal as the sole predictor of the higher frequencies. Adding predictors such as winds or their principal components derived from PCA could improve the results. Also the model in Equation (3.6) does not allow for phase or frequency modulation, which seems to be one limiting factor in the predictability of the signal. Nonetheless this preliminary step shows that it could be possible to predict seiche signals using simple physical predictors, but that the relationship between these predictors and the predicted quantity is probably nonlinear. More complex models could be tested, including machine learning techniques. Once the demonstration of a skillful model is made, such an approach can be extended in 2D using simulated (i.e., gridded) fields. In particular extensions of the VMD method exist for multivariate time series, e.g., Dynamic Mode Decomposition (DMD) [8].

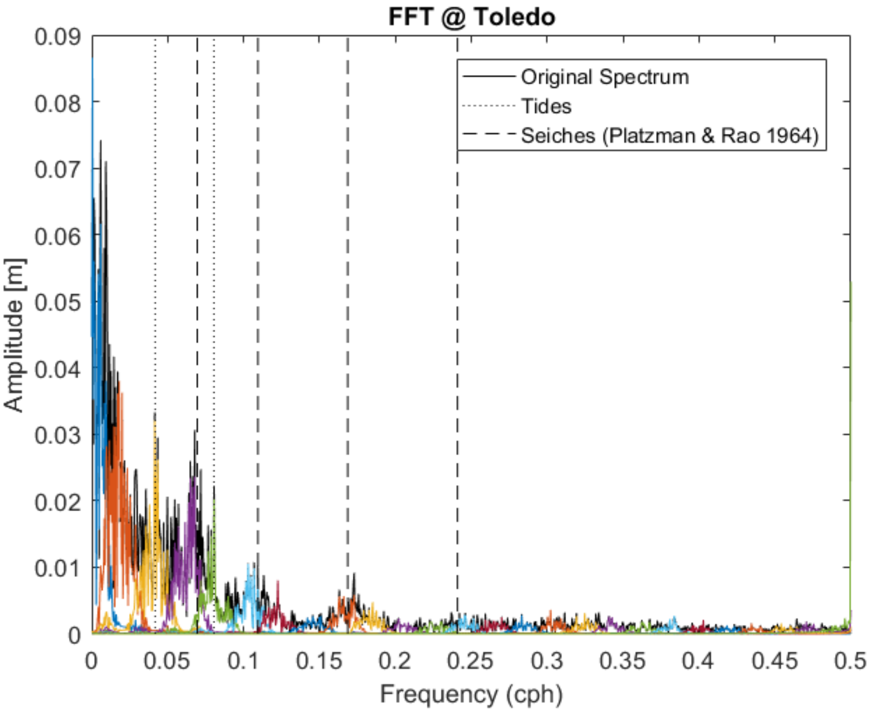

Figure 3.8: FFT amplitude spectrum for each VMD mode (colours) compared to the original spectrum (black) at Toledo (Ohio) in Lake Erie. Vertical lines denote major diurnal and semi-diurnal tidal frequencies (dotted) and documented seiche frequencies (dashed).

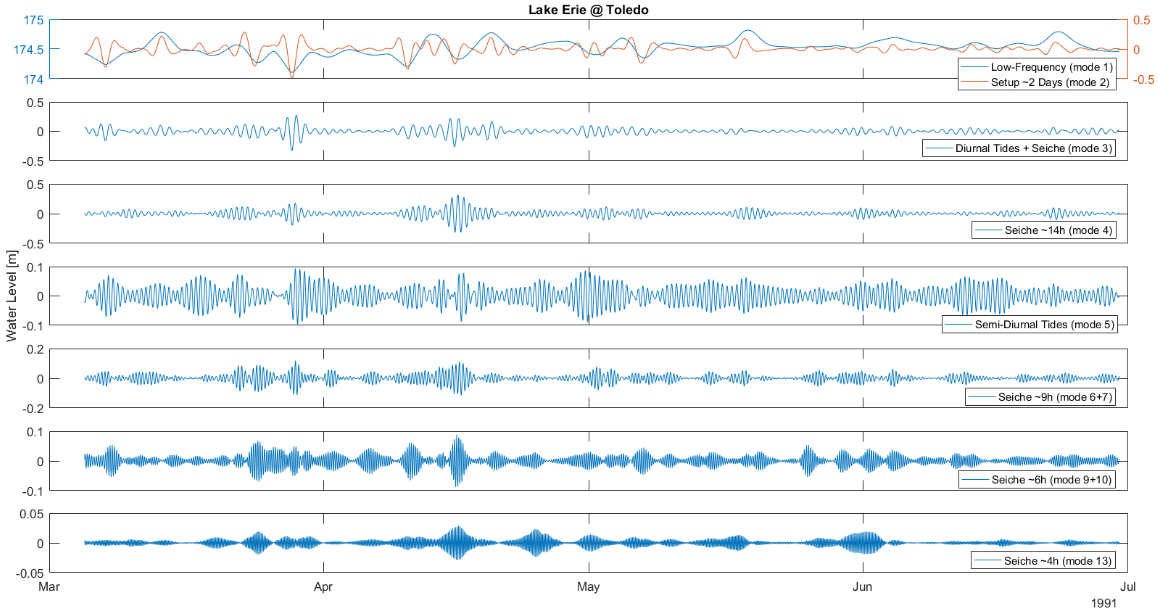

Figure 3.9: Reconstructed time series for selected groups of modes corresponding to relevant tidal and/or seiche periods.

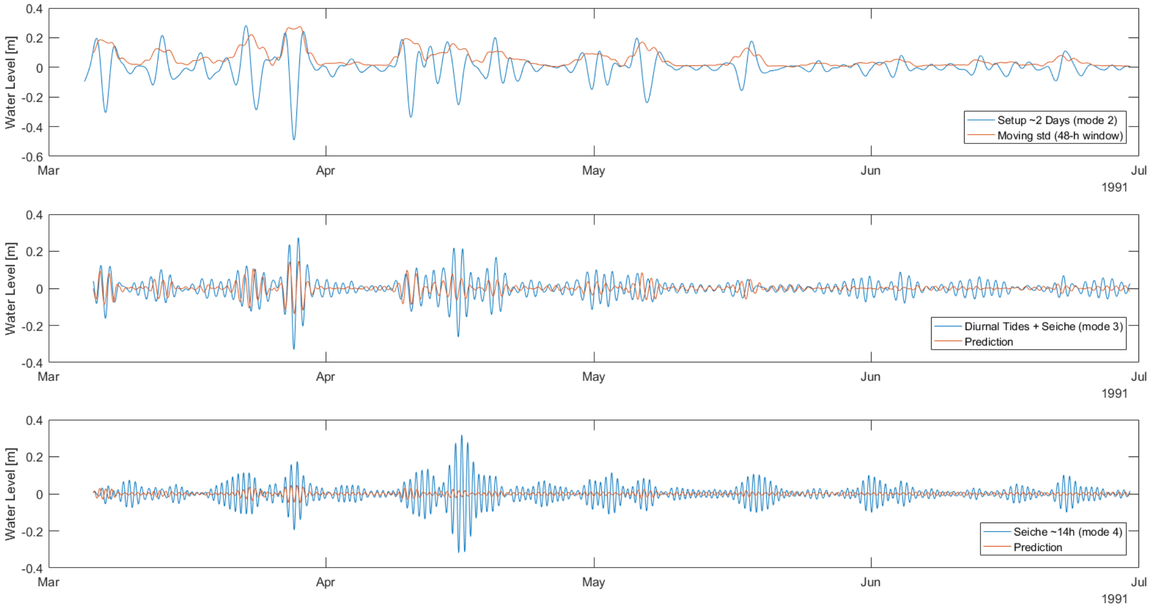

Figure 3.10: Water level predictions for the diurnal mode and 14-h seiche mode using the setup signal as the sole predictor.

3.6 A reduced model approach

In this section we focus on understanding only the seiche dynamics in Lake Erie, by excluding the physical phenomena that do not have much effect on them and thus obtaining a simpler system that can be solved without large computational resources.

The computational approach proposed in [5] uses the following model for the shallow water waves:

$$\frac{\partial h'}{\partial t'} + \nabla' \cdot \boldsymbol{q}' = 0, \tag{3.8a}$$

$$\frac{\partial \boldsymbol{q}'}{\partial t'} + \nabla' \cdot \left(\frac{\boldsymbol{q}'\boldsymbol{q}'^T}{H'} \right) + c'^2 \nabla' h' - \frac{1}{\rho'} \nabla' \cdot (H' \tau') + \frac{\boldsymbol{\tau}'_b}{\rho'} - f'_c \left(\boldsymbol{q}' \times \boldsymbol{e}_z \right) = \frac{\boldsymbol{\tau}'_s}{\rho'}, \tag{3.8b}$$

where the dimensional unknowns $\boldsymbol{q}'$ and h' are the specific discharge and water level perturbation, $H' = h' - z'$ is the water level with (known) lake depth z', $c'^2 = gH'$ is the celerity of waves, τ' is the Reynolds stress, and $\boldsymbol{\tau}'_{s,b}$ are the surface and bottom frictions respectively. Various terms in Equation (3.8b) correspond to different physical laws present in the model, namely advection, gravity, turbulent viscosity, bottom friction, Coriolis force, and surface friction. As seiches are a fairly simple type of water level dynamics, we expect to be able to describe them with a simpler equation than the system Equation (3.8) above. In other words, we expect some of the terms in Equation (3.8b) to be small in the seiche regime, so we should be able to drop them without losing much of the model accuracy. To understand which terms are small, we will nondimensionalize the system: both the dependent and independent variables will be scaled, which will allows us to “factor out” the typical size of each term and compare the terms with one another.

3.6.1 Nondimensionalization

To nondimensionalize system (3.8) we will write each dimensional variable as a

$$\text{dimensional variable'} = \underset{(known)}{\text{typical size}^*} \times \underset{\mathcal{O}(1)}{\text{dimensionless variable}},$$

which will allows us to see the relative sizes of different terms. Specifically, we will use the following reference quantities.

- $L^* = 389$ km: maximum length of lake Erie; set $x' = L^* x$ and $y' = L^* y$.
- $z^* = 19$ m: average depth of lake Erie; set $z' = z^* z$ and $H' = z^* H$.
- $u^* = 1$ m/s: typical water velocity; set $q' = u^* z^* q$.
- $h^* = 1$ m: typical water level perturbation; set $h' = h^* h$.
- $t^* = \frac{L^*}{u^*}$ s: time scale (chosen for consistency).

For clarity we will first nondimensionalize the stress terms, beginning with the Reynolds stress:

$$\tau' = \rho_{\text{water}} \left(\nu_l' + \nu_t'\right) \begin{bmatrix} 2\frac{\partial u_x'}{\partial x'} & \frac{\partial u_x'}{\partial y'} + \frac{\partial u_y'}{\partial x'} \\ \frac{\partial u_x'}{\partial y'} + \frac{\partial u_y'}{\partial x'} & 2\frac{\partial u_y'}{\partial y'} \end{bmatrix}$$

where laminar viscosity ν_l' equals 10^{-3} m^2/s and turbulent viscosity is given by the following equation.

$$\nu_t' = H'^2 \sqrt{2\left(\frac{\partial u_x'}{\partial x'}\right)^2 + 2\left(\frac{\partial u_y'}{\partial y'}\right)^2 + \left(\frac{\partial u_x'}{\partial y'} + \frac{\partial u_y'}{\partial x'}\right)^2} \tag{3.9}$$

$$= \underbrace{\frac{z^{*2} u^*}{L^*}}_{=:\nu^*} \underbrace{H^2 \sqrt{2\left(\frac{\partial u_x}{\partial x'}\right)^2 + 2\left(\frac{\partial u_y}{\partial y'}\right)^2 + \left(\frac{\partial u_x}{\partial y'} + \frac{\partial u_y}{\partial x'}\right)^2}}_{=:\nu_t} \tag{3.10}$$

Nondimensionalizing the Reynolds stress can be carried out as follows.

$$\tau' = \rho_w \left(\nu_l' + \nu_t'\right) \begin{bmatrix} 2\frac{\partial u_x'}{\partial x'} & \frac{\partial u_x'}{\partial y'} + \frac{\partial u_y'}{\partial x'} \\ \frac{\partial u_x'}{\partial y'} + \frac{\partial u_y'}{\partial x'} & 2\frac{\partial u_y'}{\partial y'} \end{bmatrix} = \underbrace{\frac{\rho_w \nu^* u^*}{L^*}}_{=:\tau^*} \underbrace{\left(\frac{\nu_l'}{\nu^*} + \nu_t\right)}_{\approx 1} \underbrace{\begin{bmatrix} 2\frac{\partial u_x}{\partial x} & \frac{\partial u_x}{\partial y} + \frac{\partial u_y}{\partial x} \\ \frac{\partial u_x}{\partial y} + \frac{\partial u_y}{\partial x} & 2\frac{\partial u_y}{\partial y} \end{bmatrix}}_{=:\tau}, \quad \tau^* = \rho_w \left(\frac{z^* u^*}{L^*}\right)^2$$

We also factor out the scaling factors from bottom and surface frictions.

$$\boldsymbol{\tau}_b' = \left(\frac{\rho_w g n'^2 |q'|}{H'^{7/3}}\right) \boldsymbol{q}' = \underbrace{\frac{\rho_w g n'^2 u^{*2}}{z^{*1/3}}}_{=:\tau_b^*} \underbrace{\left(\frac{|q|}{H^{7/3}}\right) \boldsymbol{q}}_{=:\boldsymbol{\tau}_b}$$

$$\boldsymbol{\tau}_s' = \rho_a C_w' |\boldsymbol{w}'| \boldsymbol{w}' = \underbrace{\rho_a C_w' w^{*2}}_{=:\tau_s^*} \underbrace{|\boldsymbol{w}| \boldsymbol{w}}_{=:\boldsymbol{\tau}_s}$$

Note that n' is the Manning coefficient and C_w is the wind drag coefficient.

We are now ready to nondimensionalize the original system (3.8).

$$\frac{\partial h}{\partial t} + \frac{z^*}{h^*} \nabla \cdot \boldsymbol{q} = 0,$$

$$\frac{\partial \boldsymbol{q}}{\partial t} + \nabla \cdot \left(\frac{\boldsymbol{q}\boldsymbol{q}^T}{H}\right) + \frac{g h^*}{u^{*2}} H \nabla h - \frac{t^* z^{*2} u^*}{L^{*3}} \nabla \cdot (H\tau) + \frac{t^* g' n'^2 u^*}{z^{*4/3}} \boldsymbol{\tau}^b - t^* f_c' \left(\boldsymbol{q} \times \boldsymbol{e}_z\right) = \frac{\rho_a C_w w^* t^*}{\rho_w u^* z^*} \boldsymbol{\tau}^s$$

Substituting typical parameter values into the above system yields the following.

$$\frac{\partial h}{\partial t} + \{10\} \nabla \cdot \boldsymbol{q} = 0,$$

$$\frac{\partial \boldsymbol{q}}{\partial t} + \{1\} \nabla \cdot \left(\frac{\boldsymbol{q}\boldsymbol{q}^T}{H}\right) + \{10\} H \nabla h - \{10^{-8}\} \nabla \cdot (H\tau) + \{74\} \boldsymbol{\tau}^b - \{6.35\} \left(\boldsymbol{q} \times \boldsymbol{e}_z\right) = \{1.2 \times 10^{-2} w^*\} \boldsymbol{\tau}^s$$

Clearly the Reynolds stress term can be dropped, which already simplifies the system noticeably. It is still nonlinear, however, and therefore difficult to work with. To continue our analysis we will concentrate on a particular scenario.

3.6.2 Linearization of the seiche scenario

To study the seiche oscillation scenario we assume that a unidirectional wind has created a setup of the form

$$h'(x',y') \approx 2h^* \left(\frac{x'}{L^*} - \frac{1}{2}\right), \quad h^* \ll z^*$$

and then stopped. Then after dropping the Reynolds stress term we are left with

$$\frac{\partial h}{\partial t} + \frac{1}{\epsilon}\nabla \cdot \boldsymbol{q} = 0, \tag{3.11a}$$

$$\frac{\partial \boldsymbol{q}}{\partial t} + \nabla \cdot \left(\frac{\boldsymbol{q}\boldsymbol{q}^T}{H}\right) + \alpha\epsilon H \nabla h + \frac{t^* g n'^2 u^*}{z^{*4/3}}\boldsymbol{\tau}^b - t^* f_c' \left(\boldsymbol{q} \times \boldsymbol{e}_z\right) = 0 \tag{3.11b}$$

where $\epsilon = \frac{h^*}{z^*}$ is assumed to be small and α equals $\frac{gz^*}{u^{*2}}$. Suppose h has a regular expansion in powers of ϵ:

$$h = h_0 + \epsilon h_1 + \mathcal{O}(\epsilon^2). \tag{3.12}$$

Then to balance the two terms in Equation (3.11a) the expansion for q must start at first order:

$$q = \epsilon q_1 + \mathcal{O}(\epsilon^2). \tag{3.13}$$

We can now linearize system (3.11) by plugging in representations (3.12) and (3.13), and only keeping leading-order terms in each equation. The advection and bottom stress terms are both quadratic in $\boldsymbol{q}$, so they do not appear in the leading order system. Moreover we have

$$H\nabla h = (\epsilon h - z)\nabla h = -z\nabla h + \mathcal{O}(\epsilon),$$

which will linearize the gravity term. The linearized system is

$$\frac{\partial h_0}{\partial t} + \nabla \cdot \boldsymbol{q}_1 = 0, \tag{3.14a}$$

$$\frac{\partial \boldsymbol{q}_1}{\partial t} - \alpha z \nabla h_0 - t^* f_c' \left(\boldsymbol{q}_1 \times \boldsymbol{e}_z\right) = 0. \tag{3.14b}$$

We now observe that $\alpha \approx 200 \gg t^* f_c' \approx 6.3$, so the Coriolis term can be dropped too. Next we take the time derivative of (3.14a) and divergence of (3.14b) and combine the results into a Helmholtz equation for h_0:

$$\frac{\partial^2 h_0}{\partial t^2} + \alpha \nabla \cdot \left(z \nabla h_0\right) = 0.$$

To find oscillatory solutions we consider the time-harmonic Ansatz for h, i.e. we write $h(x,y,t) = e^{i\omega t} h(x,y)$. This converts the wave equation above into an eigenvalue problem for $h(x,y)$ with $\lambda = \omega^2/\alpha$:

$$\nabla \cdot \left(z \nabla h_0\right) = \lambda h, \qquad \frac{\partial h}{\partial \hat{n}}\Big|_{\partial\Omega} = 0. \tag{3.15}$$

which can be readily solved using a Finite Element Method. We also observe that the dimensional period of an oscillation is recovered by

$$T' = \frac{2\pi}{\omega'} = \frac{2\pi t^*}{\omega} = \left\{\frac{2L^*}{\sqrt{gz}}\right\}\sqrt{\frac{\pi^2}{\lambda}}, \tag{3.16}$$

where $\{\cdots\}$ is the Merian's formula for the longest natural period, while $\sqrt{\pi^2/\lambda}$ is indeed expected to be very close to unity for the dominant mode on a long lake.

3.6.3 Computational results

As a final step we used the MFEM package [6] to discretize and solve the eigenvalue problem (3.15). As the real lake is a non-convex domain with a rather rough boundary, we solve the PDE on an approximate domain: the protruding dry land containing the Long Point National Wildlife Area is ignored, and the domain itself is smoothed with a short-time heat flow to regularize the boundary somewhat. Moreover we only deal with the case of constant depth $z' \equiv 19$m. The computational results we obtained are presented in Table 3.1 and in Figure 3.11 below. The most important result is the dominant mode: even with all the approximations made during our analysis, we still get a very good agreement between the computed and observed periods. The corresponding eigenstate gives the shape of the water level perturbation that would oscillate in time after the initial setup is complete and the wind ceases.

Table 3.1: Seiche periods, in hours.

Mode	Observed	Reduced Model	Relative Error
1	14.38	13.68	4.9%
2	9.14	7.67	16%
3	5.93	5.1	14%
4	4.15	4.07	1.9%

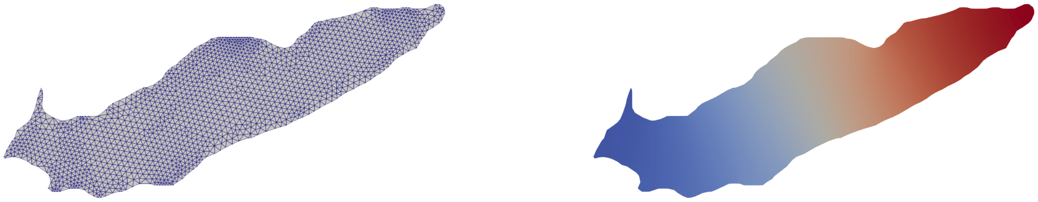

(a) FEM mesh; 1723 unknowns. **(b) Mode 1; $T_1 = 13.68$h.**

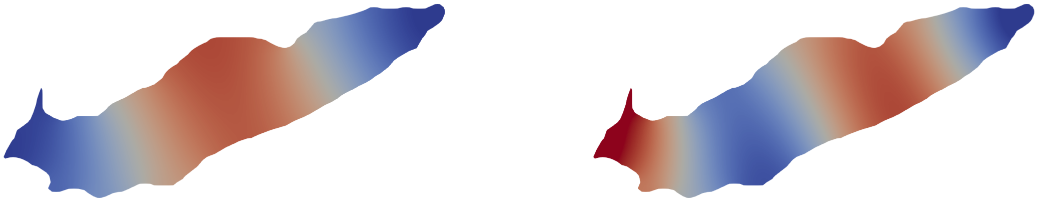

(c) Mode 2; $T_2 = 7.67$h. **(d) Mode 3; $T_3 = 5.1$h.**

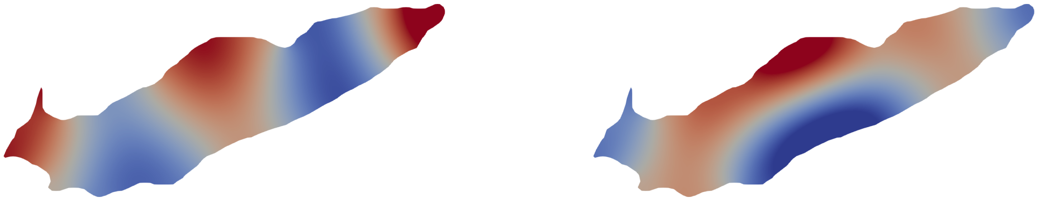

(e) Mode 4; $T_4 = 4.07$h. **(f) Mode 5; $T_5 = 3.29$h.**

Figure 3.11: Eigenmodes of Lake Erie.

The analysis presented in this section is only a crude, leading-order approximation of the seiche dynamics. The computational approximations mentioned in the previous paragraph were made largely because of time constraints and can be removed with a more careful mesh construction. The asymptotic approximations made in the previous subsection can be used to derive further corrections to both the period and the eigenstate, notably by reintroducing dissipation into the problem and estimating the number of oscillations any given state can undergo before decaying to zero. Finally further application of spectral methods [2] can be used to understand better the effect of a continuing wind.

Bibliography

[1] Konstantin Dragomiretskiy and Dominique Zosso. Variational Mode Decomposition. IEEE Transactions on Signal Processing, 62(3):531–544, February 2014. Conference Name: IEEE Transactions on Signal Processing.

[2] Mark S Gockenbach. Partial differential equations: analytical and numerical methods, volume 122. SIAM, 2005.

[3] P. F. Hamblin. Meteorological Forcing and Water Level Fluctuations on Lake Erie. Journal of Great Lakes Research, 13(4):436–453, January 1987.

[4] Mourad Heniche, Yves Secretan, Paul Boudreau, and Michel Leclerc. A two-dimensional finite element drying-wetting shallow water model for rivers and estuaries. Advances in Water Resources, 23(4):359–372, 2000.

[5] Pascal Matte, Yves Secretan, and Jean Morin. Hydrodynamic modeling of the St. Lawrence fluvial estuary. i: Model setup, calibration, and validation. Journal of Waterway, Port, Coastal, and Ocean Engineering, 143(5):04017010, 2017.

[6] MFEM: Modular finite element methods [Software]. `mfem.org`.

[7] George W. Platzman and Desiraju B. Rao. Spectra of Lake Erie water levels. Journal of Geophysical Research (1896–1977), 69(12):2525–2535, 1964.

[8] Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. Journal of Fluid Mechanics, 656:5–28, August 2010. Publisher: Cambridge University Press.

4 Load forecasting inside electrical substations on the Hydro-Québec power grid

Alexandre Blondin Massé [a]
Charlie Hébert-Pinard [b]
Alain Marcotte [c]
Olivier Milon [c]
Maksym Shpakovych [d]
Zuming Sun [e]

[a] *UQAM and IREQ*
[b] *UQAM*
[c] *Hydro-Québec*
[d] *Université de Limoges*
[e] *University of Calgary*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Massé, Hébert-Pinard, Marcotte, Milon, Shpakovych, Sun

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract: *Forecasting short-term electricity demand is a key activity for Hydro-Québec (HQ). It is necessary to support objectives as varied as the management of production or the management and maintenance of the electrical network. To properly balance supply and demand, HQ has developed expertise and internal tools to estimate the electrical load on the Québec network. An analysis of the performance of the parametric model used by HQ between 1989 and 2017, however, showed an average increase in model error since 2010. To address the problem, the "Prévision et contrôle du réseau" (PCR) department of HQ, in collaboration with the Institut de recherche d'Hydro-Québec (IREQ), has built a proof of concept showing that methods based on deep neural networks are very promising for improving significantly the accuracy of the forecast models. The goal of this workshop report is to extend those ideas to the electrical substation level.*

4.1 Introduction

At Hydro-Québec TransÉnergie et Équipement (HQTÉ), the load forecast (or "MW forecast") in operations currently relies on parametric models consisting of nonlinear modelling functions [5]. These models are fed by weather observations and forecasts, the latter of which can be corrected in a semi-automated way according to the trend observed in the last hours. Although these parametric models have been found useful and remain accurate most of the time, a slow degradation of their ability to anticipate Québec's needs has been observed. In parallel the need for precision has increased steadily because of the exportation of electricity to the United States. This degradation seems to be linked to several phenomena, each of which may have its own time scale: changes in social behaviour, in the geographical distribution of the population, in home automation equipment, and in the comfiguration of electro-intensive industry. Climate change also plays a role, in particular because of the increased needs in air conditioning during the summer.

This degradation — especially noticeable during certain periods, particularly winter — results in cumulative losses amounting to several million dollars, induced either by short-term purchases of energy on the outside networks, by the loss of sales on these same networks, or by the adoption of costly management means to make the Québec network more reliable — such as the use of gas turbines. Some means have an environmental impact because these management methods are non-renewable and emit greenhouse gases. More recently, proofs of concept developed and operationalized by the Institut de recherche d'Hydro-Québec (IREQ) and the HQTÉ "Prévision et contrôle du réseau" (PCR) unit have shown that the use of neural networks allows one to build new forecasting models that can adapt quickly to network behaviour while maintaining the provision of excellent quality forecasts. Nonetheless neural models, like parametric models, are dependent on the quality of inputs, meteorological observations, and forecasts. Also, in the case of neural models, the "physical" interpretation of the impact of each input variable on the Québec consumption is lost, because of the black-box nature of the neural approach.

The short-term load forecast problem has been extensively studied in the literature. A meta-analysis of 240 papers published between 2000 and 2019 on forecasting short-term electricity consumption [6] indicates that by far the most popular class of methods is the use of artificial neural networks, with 21 % of the proposed models making exclusive use of them, a proportion that rises to 45 % when we also count the hybrid models that include neural networks. The following models were also proposed: classical time series models, including autoregressive moving average (ARMA) models (10 % of the proposed models), regression models (9 %), and fuzzy logic and support vector regression (SVR) models (each representing around 4 %). Other approaches are based on particle swarm optimization (PSO), Bayesian vector autoregression (BVAR), decomposition models, Kalman filters, self-adaptive maps or Kohonen maps (also called self-organizing maps or SOM), Grey Prediction, ant colony algorithms (ACO), and genetic algorithms (GA). Overall machine learning techniques were used in about 43 % of the articles, hybrid techniques employing multiple models in 44 %, and statistical methods in about 13 %.

Within the family of ARMA models, a specific approach called multiple equation time series has shown promising results. In such models each period or time step of the day is treated as a separate forecasting problem, with its own equation [6]. This approach has the potential of achieving very competitive prediction accuracy [4, 8]. Initiated by Ramanathan *et al.* in 1997 [7], the approach has been significantly improved since then. For instance, in 2003, Cottet and Smith proposed an hybrid model based on a multiple equation time series together with a Bayesian approach in a case study of the regional market of New South Wales [3]. In a similar fashion, Cancelo, Espasa, and Grafe exploited the dynamism of the doubly seasonal ARIMA models for the Spanish electricity market [1]. More recently, Clements, Hurn, and Li improved the way to model the interaction between weekly and annual seasonality patterns by recognizing the importance of intraday correlation in forecasting Australian consumption [2]. The advantages of the multi-equation approach are that the explanatory factors that determine forecast performance are visible, testable, and interpretable and that the specification of model parameters is linear, which means that the ordinary least squares method can be used to estimate the parameters, rather than a numerical optimization algorithm [2].

4.2 Problem description

Building upon the classical and more recently developed models, HQ would like to transfer the forecasting process to a more granular scale with respect to the spatial dimension. Many scales can be considered: the whole province scale (currently), the regional or zonal scale, the electrical substation scale, or even the source, using data collected from smart meters. Because of several functional constraints, the electrical substation level seems to be a good starting point. Figure 4.1 illustrates the components and the layout of a generic electrical substation. In particular substations bridging the transport and the distribution layers, called “satellite substations,” are interesting candidates. HQ operates around 300 such satellite substations.

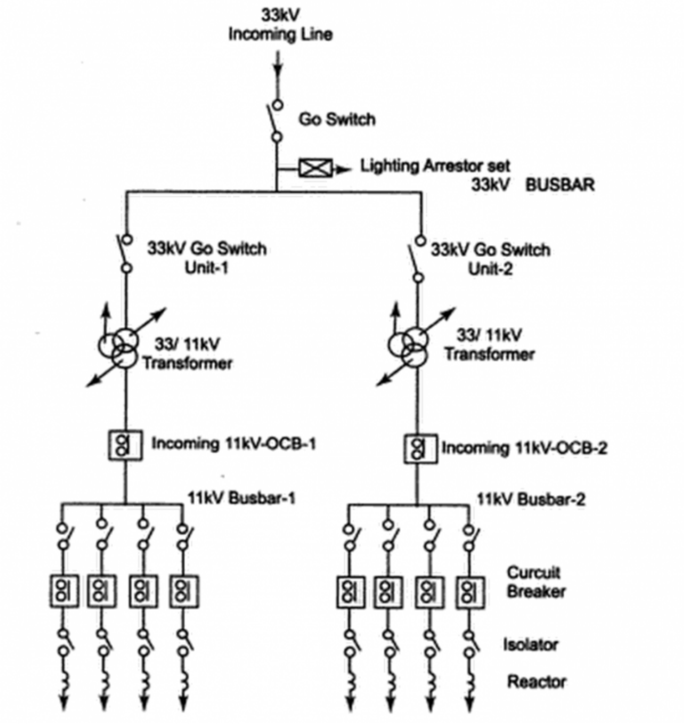

Figure 4.1: A single line diagram (SLD) of a 33kV electrical substation (source: WatElectrical.com).

In addition to forecasting the MW at the electrical substation, HQ would like to adapt the models so that they can estimate other physical quantities going through different components.

1. **Transformers**. The active power (in MW) and the reactive power (in MVar) for each transformer.
2. **Bars**. The current (in A) for each bar.
3. **Lines**. The current (in A) for each of the 3 phases.

There are several difficulties to be addressed. Clearly there are strong dependencies between the different physical quantities mentioned above, but these dependencies are hard to characterize formally, since they depend on the layout (or topology) of the substation, which varies significantly between substations. There are also multiple possible substation configurations or states.

In order to study the problem, a dataset was assembled by HQ as follows.

1. The dataset consists of 3 electrical substations.
2. All given time series span the period between January 1st, 2019 and August 3rd, 2022, at an hourly time step.
3. The load consumption for the whole Québec province, called "besoin québécois", for each time step, is also provided.
4. For each transformer of each substation, the dataset includes the active and reactive powers (respectively in MW and in MVar) passing through the transformer at each time step.
5. For each bar of each substation, the dataset includes the current (in A) going through it at each time step.
6. For each line of each substation, the dataset includes the current (in A) going through it at each time step and for each of three phases.
7. The dataset includes the configuration setting of each substation, i.e., the state of each breaker (open or closed) at each time step.
8. The dataset includes the weather conditions (temperature, wind speed, cloud coverage, and precipitation type) measured in Montréal at each time step.

A screenshot of one table is depicted in Figure 4.2.

4.3 Adopted methodology

The team assembled to tackle this problem consisted of the following persons.

- **Alain Marcotte** and **Olivier Milon**, both working in the "Prévision et contrôle du réseau" (PCR) team at Hydro-Québec.
- **Charlie Hébert-Pinard** (Université du Québec à Montréal), **Maksym Shpakovych** (Université de Limoges), and **Zuming Sun** (University of Calgary) are Master's and Ph.D. candidates interested in working on the problem submitted by HQ.
- **Alexandre Blondin Massé** (Institut de recherche d'Hydro-Québec and Université du Québec à Montréal) was in charge of coordinating the activities of the group.

The activities were scheduled as follows during the week:

Monday: Presentation of the problem, gathering of the team in the designated room, questions and answers about the domain with the HQ representatives, retrieval of the dataset;

Tuesday: Exploration of the dataset, preparation and cleaning of the dataset, description of some models used currently at Hydro-Québec;

Wednesday: Implementation of various models (based on statistical and machine-learning methods), reports on preliminary results;

Thursday: Improvement of the models, additional reports, preparation of the final presentation.

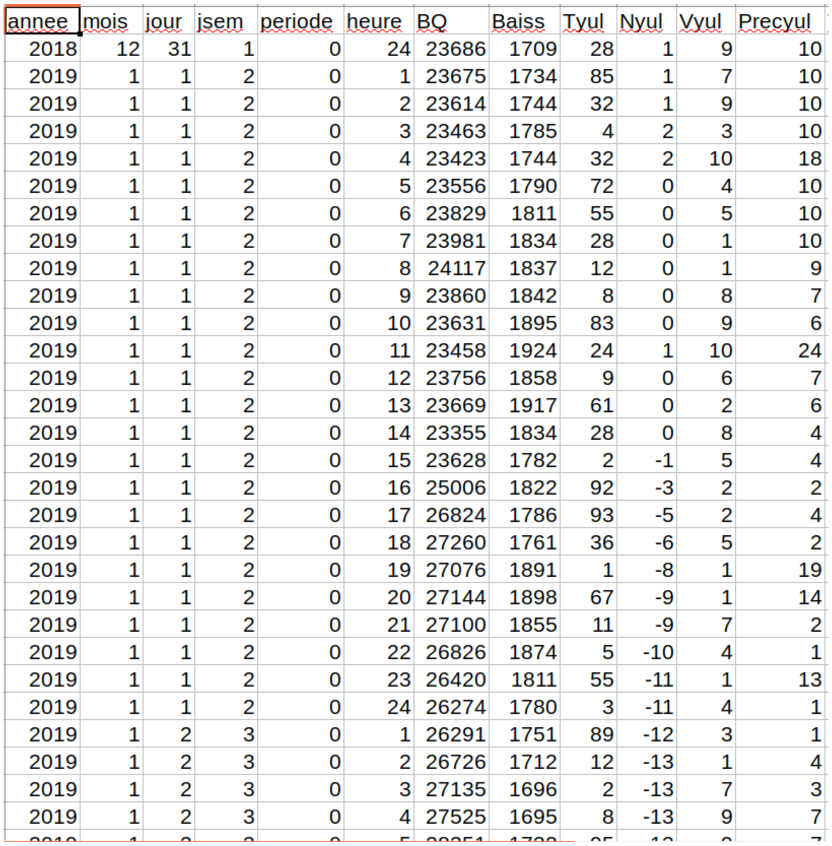

annee	mois	jour	jsem	periode	heure	BQ	Baiss	Tyul	Nyul	Vyul	Precyul
2018	12	31	1	0	24	23686	1709	28	1	9	10
2019	1	1	2	0	1	23675	1734	85	1	7	10
2019	1	1	2	0	2	23614	1744	32	1	9	10
2019	1	1	2	0	3	23463	1785	4	2	3	10
2019	1	1	2	0	4	23423	1744	32	2	10	18
2019	1	1	2	0	5	23556	1790	72	0	4	10
2019	1	1	2	0	6	23829	1811	55	0	5	10
2019	1	1	2	0	7	23981	1834	28	0	1	10
2019	1	1	2	0	8	24117	1837	12	0	1	9
2019	1	1	2	0	9	23860	1842	8	0	8	7
2019	1	1	2	0	10	23631	1895	83	0	9	6
2019	1	1	2	0	11	23458	1924	24	1	10	24
2019	1	1	2	0	12	23756	1858	9	0	6	7
2019	1	1	2	0	13	23669	1917	61	0	2	6
2019	1	1	2	0	14	23355	1834	28	0	8	4
2019	1	1	2	0	15	23628	1782	2	-1	5	4
2019	1	1	2	0	16	25006	1822	92	-3	2	2
2019	1	1	2	0	17	26824	1786	93	-5	2	4
2019	1	1	2	0	18	27260	1761	36	-6	5	2
2019	1	1	2	0	19	27076	1891	1	-8	1	19
2019	1	1	2	0	20	27144	1898	67	-9	1	14
2019	1	1	2	0	21	27100	1855	11	-9	7	2
2019	1	1	2	0	22	26826	1874	5	-10	4	1
2019	1	1	2	0	23	26420	1811	55	-11	1	13
2019	1	1	2	0	24	26274	1780	3	-11	4	1
2019	1	2	3	0	1	26291	1751	89	-12	3	1
2019	1	2	3	0	2	26726	1712	12	-13	1	4
2019	1	2	3	0	3	27135	1696	2	-13	7	3
2019	1	2	3	0	4	27525	1695	8	-13	9	7

Figure 4.2: Screenshot of one table belonging to the sample dataset.

The sample data presented some anomalies that had to be fixed: the measurements of some physical quantities expressed in MW and in A were a little below 0, which does not make sense for the current problem, so they were clamped to 0.

Given a vector of n observed values Y and a vector of n forecasted values $\hat{Y}$ for Y, recall that the mean squared error (MSE) of the forecast is given by

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}\left(Y_i - \hat{Y}_i\right)^2, \tag{4.1}$$

while the mean absolute percentage error (MAPE) is given by

$$\text{MSE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{Y_i - \hat{Y}_i}{Y_i}\right|. \tag{4.2}$$

Since the MAPE computation can be problematic around zero or small positive values, we decided to use the MSE as a reference metric when evaluating the performance of the forecasting models.

We now focus on two models that were explored in more detail during the workshop.

4.4 Evaluated models

Since the HTQÉ parametric model is not easily accessible, we used as a base reference a multiple equation model developed by IREQ, whose performance is more or less comparable.

Let $\bar{t}$ be the number of time steps per day ($\bar{t} = 24$ for the sample provided). Also, for any time step t of day d, let

- $L_{d,t}$ be the logarithm of the load (in MW) at time step t of day d;
- $T_{d,t}$ be the temperature (in Celsius) at time step t of day d;
- $W_{d,p}$ be a binary variable indicating whether d is the p-th day of the week;
- $S_{d,k}$ be a binary variable indicating whether d is a special day of kind $k \in \{1, 2, \dots, k'\}$, where k' is the number of special day types.

Moreover, for any day d and time step t of the day, let

$$\begin{aligned}
Y_{d,t,q} &= 2q\pi(d\bar{t} + t)/(365.2425\bar{t}),\\
H_{d,t,1} &= \textsc{ClRamp}^{-}_{-23,13}(T_{d,t}),\\
H_{d,t,2} &= \textsc{ClRamp}^{-}_{-23,1}(T_{d,t}),\\
C_{d,t,1} &= \textsc{ClRamp}^{+}_{21,33}(T_{d,t}),\\
C_{d,t,2} &= \textsc{ClRamp}^{+}_{28,33}(T_{d,t}),
\end{aligned}$$

where

$$\begin{aligned}
\textsc{ClRamp}^{-}_{x_0,x_1}(x) &= \begin{cases} x_1 - x_0 & \text{if } x \le x_0;\\ x_1 - x & \text{if } x_0 < x \le x_1;\\ 0 & \text{if } x_1 < x; \end{cases}\\
\textsc{ClRamp}^{+}_{x_0,x_1}(x) &= \begin{cases} 0 & \text{if } x \le x_0;\\ x - x_0 & \text{if } x_0 < x \le x_1;\\ x_1 - x_0 & \text{if } x_1 < x. \end{cases}
\end{aligned}$$

Finally, for any real numbers α, β and any random variable X, let

$$\mathrm{SC}(\alpha, \beta, X) = \alpha \sin X + \beta \cos X$$

and let $\mathcal{O} = \{0, -1, \dots\}$ be the set of days offsets when considering temperature and special days (0 is for the current day, -1 for the previous day, -7 would be for the previous week).

Then the equations defining the model are given by

$$\begin{aligned}
L_{d,t} = {} & \lambda_{t,0} + \left(\sum_{p=1}^{7} \omega_{t,p} W_p\right) L_{d-1,t} + \left[\lambda_{t,1} + \sum_{q=1}^{4} \mathrm{SC}(\gamma_{t,q,1}, \gamma_{t,q,2}, Y_{d,t,q})\right] L_{d-7,t}\\
& + \lambda_{t,2} L_{d-1,\bar{t}-1} + \lambda_{t,3} L_{d,t-1}\\
& + \phi_{t,1}\varepsilon_{d-1,t} + \phi_{t,2}\varepsilon_{d-7,t} + \varepsilon_{d,t}\\
& + \sum_{o\in\mathcal{O}} \sum_{\ell\in\mathcal{L}} \sum_{k=1}^{2} \tau^{\ell}_{t,o,k,1} H^{\ell}_{d+o,t,k} + \sum_{o\in\mathcal{O}} \sum_{\ell\in\mathcal{L}} \sum_{k=1}^{2} \tau^{\ell}_{t,o,k,2} C^{\ell}_{d+o,t,k}\\
& + \sum_{o\in\mathcal{O}} \sum_{\ell\in\mathcal{L}} \sum_{k=1}^{3} \nu^{\ell}_{t,o,k} \overline{N}^{\ell}_{d+o,t,k}\\
& + \sum_{o\in\mathcal{O}} \sum_{\ell\in\mathcal{L}} \omega^{\ell}_{t,o} \overline{W}^{\ell}_{d+o,t}\\
& + \sum_{o\in\mathcal{O}} \sum_{k=1}^{k'} \sigma_{t,o,k} S_{d+o,k} + \sum_{o\in\mathcal{O}} \delta_o \mathrm{DST}_{d+o},
\end{aligned} \tag{4.3}$$

where, for each $t = 1, 2, \dots, \bar{t}$ and each $o \in \mathcal{O}$, $\lambda_{t,i}$ (for $i = 0, 1, 2, 3$), $\phi_{t,i}$ (for $i = 1, 2,$), $\omega_{t,p}$ (for $p = 1, 2, \dots, 7$), $\gamma_{t,q,i}$ (for $q = 1, 2, 3, 4$ and $i = 1, 2$), $\tau_{t,o,k,i}$ (for $k = 1, 2$ and $i = 1, 2, 3, 4$) and $\sigma_{t,o,k,i}$

(for $k = 1, 2, \dots, k'$ and $i = 1, 2$), for $t = 1, 2, \dots, \bar{t}$, $o \in \mathcal{O}$ and $\ell \in \mathcal{L}$, $\tau^{\ell}_{t,o,k,i}$ (for $k = 1, 2$ and $i = 1, 2$), $\nu^{\ell}_{t,o,k}$ (for $k = 1, 2$) and $\omega^{\ell}_{t,o}$ are the parameters to learn and $\varepsilon_{d,t}$ and $\varepsilon_{d-1,t}$ are the unexpected changes in load on the previous day and week, and $\varepsilon_{d,t}$ is the expected error of the forecast model.

The goal of this model is to predict $L_{d,t}$ using $T_{\hat{d},\hat{t}}$, $W_{\hat{d},\hat{p}}$, $S_{\hat{d},k}$ for $\hat{d} \in \mathcal{D}$, $\hat{t} \in \mathcal{T}$, and $\hat{d} \in \mathcal{P}$ where

1. $\mathcal{D}$ is the set of days used for the prediction for the selected day $d \in \mathbb{N}$;
2. $\mathcal{T}$ is the set of hours used for the prediction for the selected hour $t \in \{1, \dots, 24\}$;
3. $\mathcal{P}$ is the set of days of week used for the prediction for the selected day of week $p \in \{1, \dots, 7\}$.

$\mathcal{D}$, $\mathcal{T}$ and $\mathcal{P}$ are the historical datasets that contain the number of days, hours, weeks of the day that we use to predict the value of $L_{d,t}$. For example, we have found that information from the previous 8 hours, the previous day, and the same day of the week of the previous week is informative for the prediction.

The model described above being linear by definition, one might wonder whether a nonlinear model can capture more relations between the variables. Therefore we considered a nonlinear prediction model of the form

$$\mathrm{NN}(x) = w_2^{\top} g(W_1 x + b_1, \theta) + b_2, \tag{4.4}$$

where $W_1 \in \mathbb{R}^{m \times n}$, $w_2 \in \mathbb{R}^m$, $b_1 \in \mathbb{R}^m$, $b_2 \in \mathbb{R}$, $\theta \in \mathbb{R}^m$ are trainable parameters and $g(x, \theta)$ is an m-dimensional parametric ReLU function. Then to evaluate model prediction for the selected day d and hour t the input vector x is defined as

$$x = \left(\{L_{\hat{d},\hat{t}}\}_{\hat{d} \in \mathcal{D}, \hat{t} \in \mathcal{T}}, \{T_{\hat{d},\hat{t}}\}_{\hat{d} \in \mathcal{D}, \hat{t} \in \mathcal{T}}, \{W_{\hat{d},\hat{p}}\}_{\hat{d} \in \mathcal{D}, \hat{p} \in \mathcal{P}}, \{S_{\hat{d},k}\}_{\hat{d} \in \mathcal{D}} \right) \in \mathbb{R}^n. \tag{4.5}$$

Actually, in the course of the workshop, we used more information than shown above to construct x (cloudiness, precipitation, wind speed, ...). This information can be found in Figure 4.2. We also tried to predict the exact value of the load (i.e. $\exp(L_{d,t})$) instead of its logarithm and the historical data sets $\mathcal{D}$, $\mathcal{T}$ and $\mathcal{P}$ were chosen empirically. Also the hidden layer size m was set to 64 and the parameters $W_{1,2}$, $b_{1,2}$, θ were found after training the model (4.4) on the provided dataset.

Note that we also conducted experiments with a third model, i.e., a modification of the nonlinear one, which we will not describe here (see next section).

4.5 Experiments

Using the sample provided, experiments were conducted to compare the performance of the different models under consideration. Because of time constraints we focused only on one transformer of one specific substation.

The reference multiple-equation model had to be adapted to take into account the nature of the data, in particular the fact that the load could be zero in some cases. The resulting evaluation is summarized in Figure 4.3. As is often the case in load forecasting, the forecasts are quite good during the first and last hours of the night, since the load is more stable during these periods. The forecasts are less accurate during the day, especially in the afternoon (around 3PM). For a better understanding of the results, the forecasts were also compared with the observed values with respect to the week day (Figure 4.4(a)) and the month (Figure 4.4(b)). The Mondays, Saturdays, and Sundays are typically harder to forecast, because of the effect of the special days and week-end human behavior, which exhibits more variation. As expected the error is greater during winter months, especially December and January.

The neural network approach yielded slightly better results, as depicted in Figure 4.5. Two variants were explored: one chaining the models with a simple linear regression (the blue curve in Figure 4.5(a)) and the other with the nonlinear variant (the orange curve in Figure 4.5(a)). As a bonus a weighted version is also provided, exploiting the best of the two variants (the green curve in Figure 4.5(a)). Figure 4.6 compares the real observed values with the forecasts on an arbitrary day.

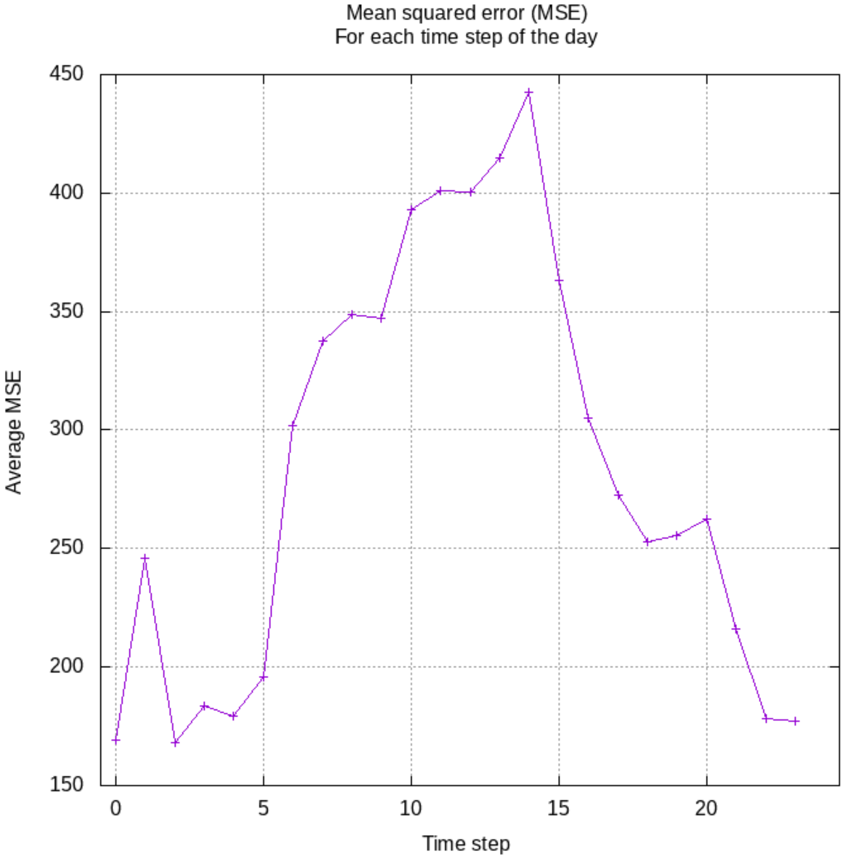

Figure 4.3: Mean squared error (MSE) averaged for each time step.

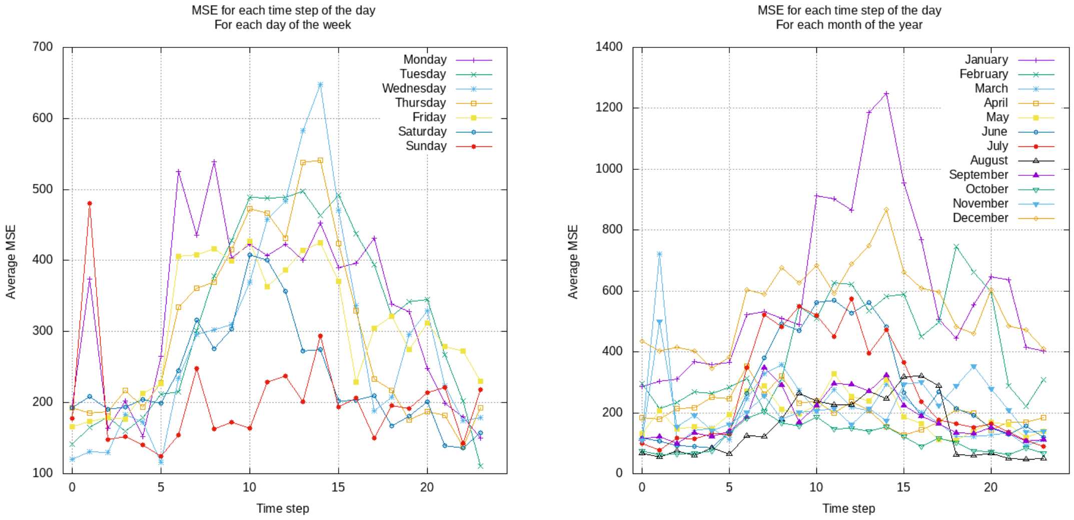

Figure 4.4: Mean squared error (MSE) averaged for each time step and (a) each week day (b) each month.

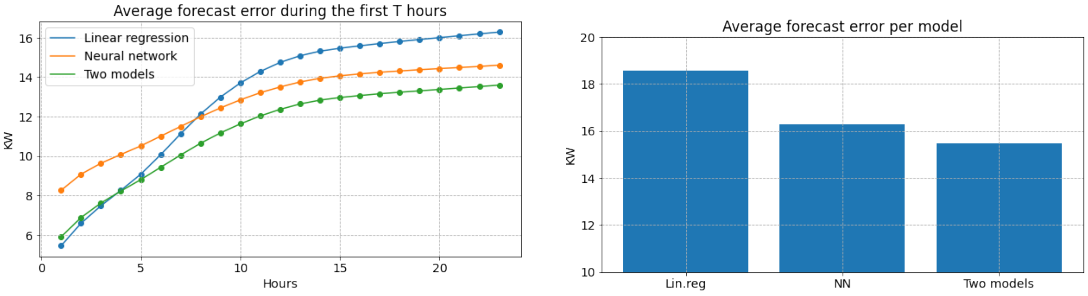

(a) (b)

Figure 4.5: Forecast quality comparison. (a) Comparison of the average forecast error between the linear and the nonlinear variants of the neural network model, and a weighted combination of them, for each of the 24 time steps of the day. (b) Comparison of the overall average forecast error of each of the three models.

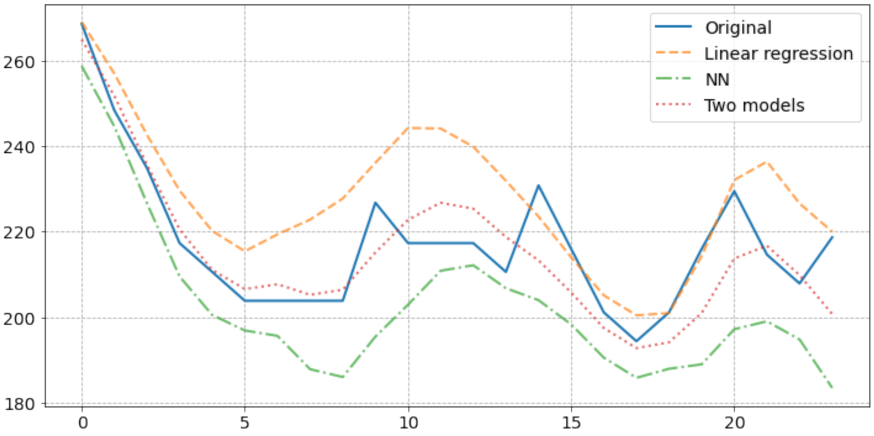

Figure 4.6: Comparison between the real power load and the load forecasted by each of the three models for an arbitrary day.

4.6 Concluding remarks

Although we did not have time to explore the submitted problem in detail, we were satisfied with the preliminary results obtained. The workshop also allowed us to identify further the difficulties of adapting the MW forecasting methods at the electrical substation level.

In the future we intend to build a global model for all elements of the system (transformers, bars, breakers), in particular using the information on the breakers' state. We also identified some strategies to improve the training, especially the introduction of parallelization. Further improvement might be obtained by using a faster optimization algorithm instead of the default stochastic gradient descent.

Bibliography

[1] J. R. Cancelo and R. Espasa, A. Grafe. Forecasting the electricity load from one day to one week ahead for the Spanish system operator. International Journal of Forecasting, 24(4):588–602, 2008.

[2] A. E. Clements, A. S. Hurn, and Z. Li. Forecasting day-ahead electricity load using a multiple equation time series approach. European Journal of Operational Research, 251(2):522–530, 2016.

[3] R. Cottet and M. Smith. Bayesian modeling and forecasting of intraday electricity load. Journal of the American Statistical Association, 98(464):839–849, 2003.

[4] M. Espinoza, C. Joye, R. Belmans, and B. De Moor. Short-term load forecasting, profile identification, and customer segmentation: a methodology based on periodic time series. IEEE Transactions on Power Systems, 20(3):1622–1630, 2005.

[5] M. Grenier. Short-term load forecasting at Hydro-Québec TransÉnergie. In 2006 IEEE Power Engineering Society General Meeting, pp. 5 pp.-, doi: 10.1109/PES.2006.1709029.

[6] A. B. Nassif, B. Soudan, M. Azzeh, I.B. Attilli, and O. AlMulla. Artificial intelligence and statistical techniques in short-term load forecasting: A review. CoRR, abs/2201.00437, 2022.

[7] R. Ramanathan, R. Engle, C. W. J. Granger, F. Vahid-Araghi, and C. Brace. Short-run forecasts of electricity loads and peaks. International Journal of Forecasting, 13(2):161–174, 1997.

[8] L. J. Soares and M. C. Medeiros. Modeling and forecasting short-term electricity load: A comparison of methods with an application to Brazilian data. International Journal of Forecasting, 24(4):630–644, 2008.

5 Turbulence in the air: Creating a heat map and building a seasonal diagram

Alberto Fornaci[a]
Adrian Ganea[a]
Brent King[a]
Katsiaryna Vashchankova[a]
Javier Almonacid[b]
Sean Bohun[c]
Douglas Bowen[d]
Slim Ibrahim[e]
Michael R Lindstrom[f]
Joy Liu[g]
Kyle Onghai[h]
Thomas Gkelsinis[i]

[a] *IATA*
[b] *Simon Fraser University*
[c] *Ontario Tech University*
[d] *Wilfrid Laurier University*
[e] *University of Victoria*
[f] *The University of Texas Rio Grande Valley*
[g] *Dalhousie University*
[h] *University of California in Los Angeles*
[i] *Université de Rouen Normandie*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Fornaci, Ganea, King, Vashchankova, Almonacid, Bohun, Bowen, Ibrahim, Lindstrom, Liu, Onghai, Gkelsinis

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

5.1 Introduction

Air transportation is an essential part of modern living with many thousands of flights taking place every day. With such popularity it is of the utmost importance to minimize risks when in flight. Despite the ubiquitous messages to "fasten your seatbelt," however, an average of 58 people in the United States are injured through turbulence events annually from failing to adhere to this messaging [1]. Airlines make use of forecasts to predict patches of turbulence. Currently these forecasts are not reliable and no remote technique can predict the intensity of turbulence. In clear air, where particulate matter is very small,[1] it is effectively invisible to weather radar. Even if turbulence were carrying particles of a sufficiently large size, radar cannot detect winds that are moving in a direction orthogonal to the direction of the radar beam. These difficulties in detecting turbulence are made more urgent by an expected increase in the global average of clear air turbulence (CAT) events.

In addition to forecasting, airlines make use of current reports from pilots when regions of turbulence are detected. These regions are identified as possibly turbulent for subsequent flights. These reports are not reliable, however, as they tend to be subjective and based on the personal experience of the pilot, which depends partly upon the aircraft mass. The International Air Transport Association (IATA) has developed a system that characterizes turbulence in real time by using the current flight data.

5.1.1 Current state

At the present time IATA is able to track the eddy dissipation rate (EDR) along the flight path by computing a real time spectral analysis of the wind velocity. This EDR measurement is correlated with the degree of turbulence experienced by an aircraft, independently of its mass, which is not the case for the subjective measurement of the rapid random acceleration of the aircraft, which depends on airplane inertia. With this technique, the wind measurement itself characterizes the intensity of the turbulence. With a platform that can sample the EDR continuously, this data can be viewed by the pilot with the help of the data visualization software. As of this writing, airlines can display EDR data over the most recent four-hour window.

In its current state, each sample of EDR is displayed as a coloured dot at its geographic location. The colour of the dot corresponds to the EDR, which is related to the intensity of the turbulence. Figure 5.1 shows a typical example of an IATA turbulence visualization dashboard. The information needed by pilots is hidden within the dots. To access it, a dot is selected and the data is displayed. This system does not allow for rapid decisions. Also the weather forecast may not be of any help, as CAT is undetectable by radar. Indeed, if pilots use only the weather forecast, they may be directed into a high CAT region.

The task proposed by IATA was to develop a new way to visualize the EDR samples in a format reminiscent of a storm tracker. Confined to the neighbourhood of the various flight paths over a four-hour window, new EDR events would be added to the map as they are acquired. After four hours, when they are considered stale, they would be removed from the map. It was thought that by animating this behaviour over longer periods of time, the underlying structure and motion of the turbulence could be tracked and predicted. More extensive models could be focused on the predicted propagation and evolution of the turbulent field, but these aspects were set aside in the present work. We focus on an attempt to characterize the properties of the observed EDR events as recorded in the IATA database and producing a sequence of images that are collected into an animation. It is expected that the animation will reveal evolving hidden turbulent structures.

---

[1] Particles of matter with diameters less than 10 μm are too small to be detected by radar systems with wavelengths on the order of centimetres.

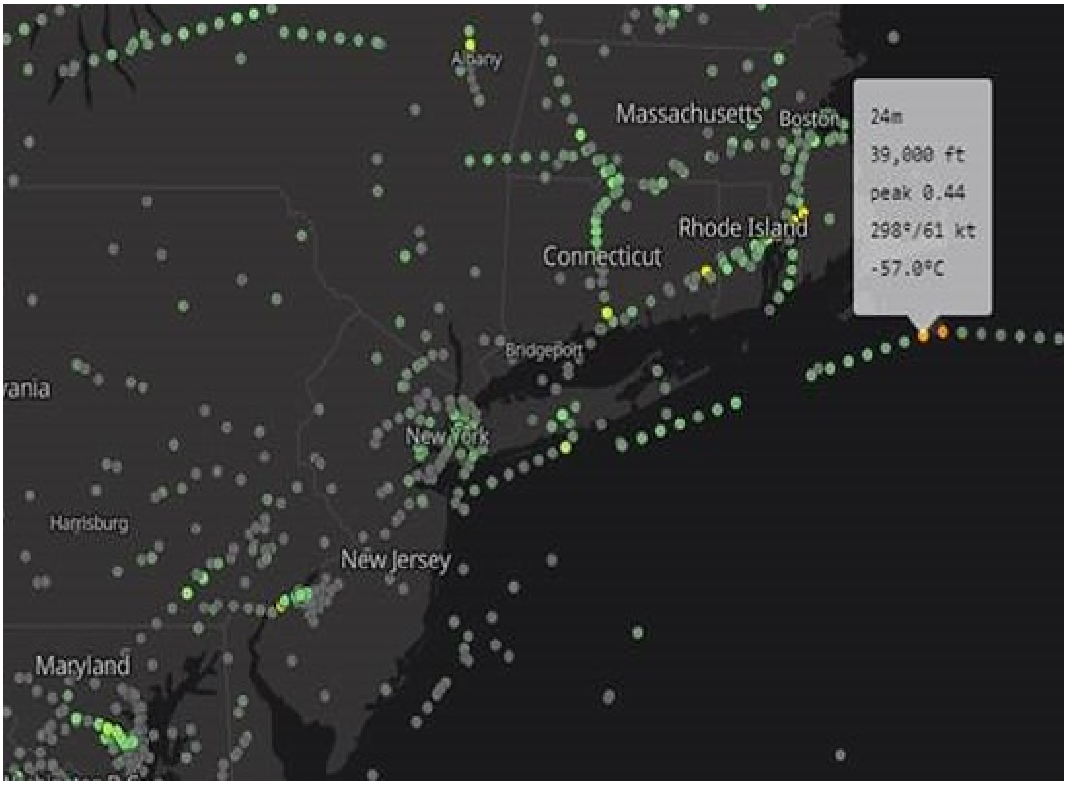

Figure 5.1: Example of the current IATA turbulence visualization.

5.2 Characterizing Turbulence

The turbulence kinetic energy (per unit mass of air) is connected to the variations in wind velocity away from its average value. In this sense the wind velocity is given by $\mathbf{u}(t) = \bar{\mathbf{u}}(t) + \mathbf{u}'(t)$, with $\bar{\mathbf{u}}(t) = (\bar{u}_1(t), \bar{u}_2(t), \bar{u}_3(t))$ denotes a running average behaviour and $\mathbf{u}' = (u_1(t), u_2(t), u_3(t))$ denotes the deviations. The energy is given by

$$E_T = \frac{1}{2}|\mathbf{u}'|^2 = \frac{1}{2}\left(u_1^2 + u_2^2 + u_3^2\right). \tag{5.1}$$

Although random in nature, this energy has a particular structure. It is distributed throughout a variety of length scales that are characterized by the eddies of motion that build up the turbulent flow. For an eddy of characteristic size l with corresponding wavenumber $k = 2\pi/l$, the energy spectrum, $E(k)\,\mathrm{d}k$, denotes the energy contained in eddies with wavenumber k to $k + \mathrm{d}k$. With respect to this description, we have

$$E_T = \int_0^\infty E(k)\,\mathrm{d}k. \tag{5.2}$$

The typical wavenumber dependence for this turbulent energy is shown in Figure 5.2.

When in equilibrium, the energy produced at large spatial scales is transferred to successively smaller scales, through an inertial subrange, to a dissipation range governed by the dynamic viscosity[2] ν. Within the inertial subrange of the spectrum, the eddy dissipation rate (EDR), ϵ, is the rate at which turbulent kinetic energy (density) is transmitted into heat. It is expressed in units of $[\epsilon] = [E_T]/[\text{time}] = (\mathsf{m}\ \mathsf{s}^{-1})^2\ \mathsf{s}^{-1} = \mathsf{m}^2\ \mathsf{s}^{-3}$. To extract the behaviour of the energy in this region we attempt to find a quantity depending only on the EDR and the wavenumber, $D = \epsilon^a k^b$, so that D has the same units

---

[2] $\nu = \mu/\rho$ where μ is the kinematic viscosity.

as $E(k)$. Since $[E(k)] = [E_T]/[k] = \mathsf{m}^3\ \mathsf{s}^{-2}$, $[k] = \mathsf{m}^{-1}$, and $[\epsilon] = \mathsf{m}^2\ \mathsf{s}^{-3}$, we have $D = \epsilon^{2/3}k^{-5/3}$. The conclusion is that if a turbulent structure is within the inertial subrange, then there is a scalar constant[3] C such that

$$E(k) = C\epsilon^{2/3}k^{-5/3}. \tag{5.3}$$

Expression (5.3) is used by the NCAR In Situ Turbulence Detection Algorithm (TDR) to characterize the level of turbulence during a flight [2]. The energy spectrum with respect to wavenumber within the inertial subrange of the observed wind speed in flight is proportional to energy dissipation rate raised to a power.

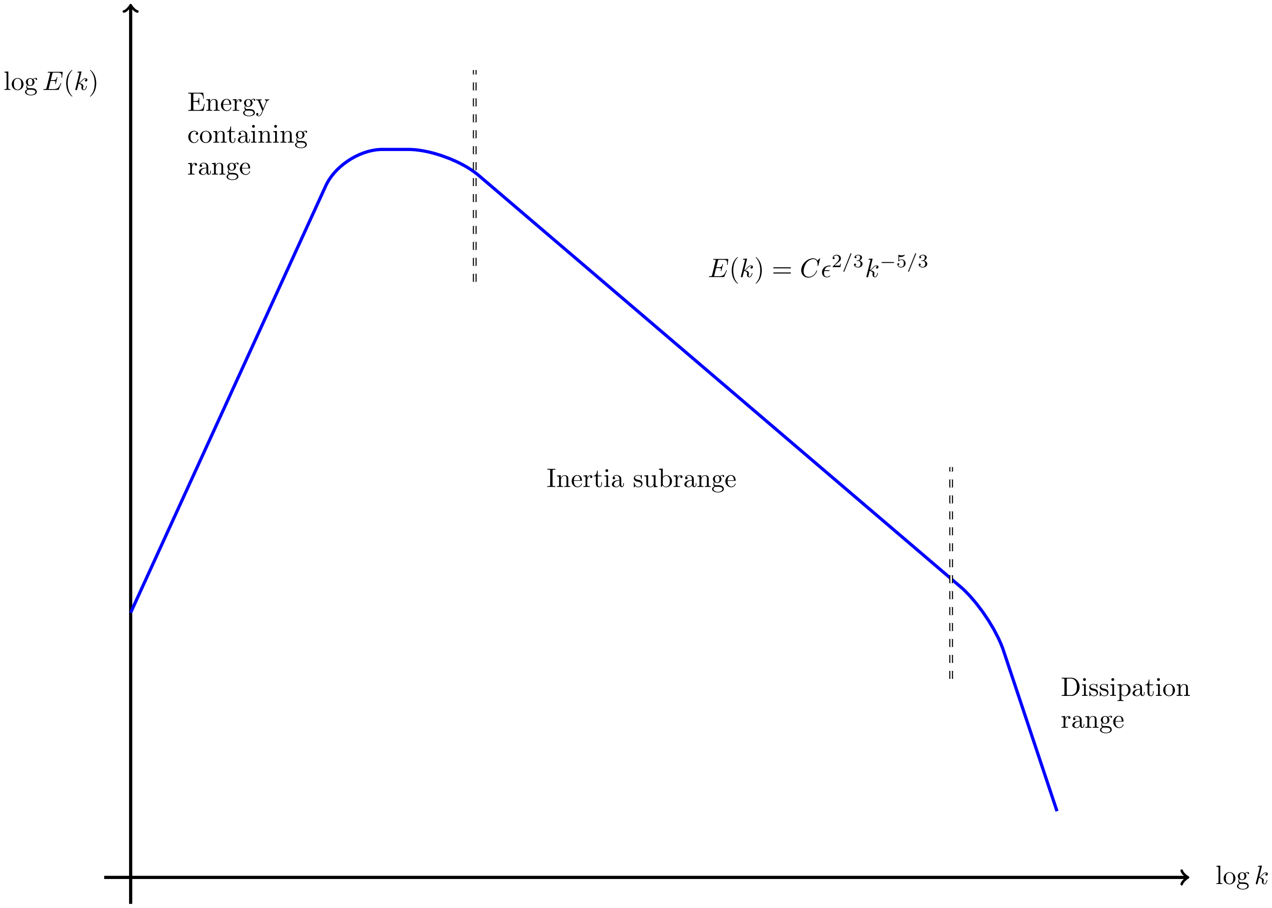

Figure 5.2: A schematic of the turbulent energy spectrum $E(k)$.

Other characteristic scales can be extracted in this region with a characteristic length, L, satisfying $k \sim \mathsf{L}^{-1}$, a characteristic time, T, with $E \sim \mathsf{L}^3\mathsf{T}^{-2}$, and associated speed $\mathsf{U} \sim \mathsf{L}\mathsf{T}^{-1}$, and Reynolds number $\mathrm{Re} \sim \mathsf{U}\mathsf{L}\nu^{-1}$. This results in

$$\mathsf{L} \sim k^{-1}, \qquad \mathsf{T} \sim \epsilon^{-1/3}k^{-2/3}, \qquad \mathsf{U} \sim \epsilon^{1/3}k^{-1/3}, \qquad \mathrm{Re} = \epsilon^{1/3}\nu^{-1}k^{-4/3}. \tag{5.4}$$

5.2.1 Real-time estimation of the EDR

For data acquisition, the acronym EDR does not refer to ϵ, but rather to $\epsilon^{1/3}$ ($[\mathrm{EDR}] = \mathsf{m}^{2/3}\ \mathsf{s}^{-1}$). The quantity $\epsilon^{1/3}$ is reported rather than ϵ, as $\epsilon^{1/3}$ can be considered as an effective standard deviation of the wind speed [3]. This measurement of the EDR is independent of the eddy size used for its

---

[3] The dimensionless constant $C \sim 1.5$ can be determined experimentally.

calculation, provided that both large and small eddies are avoided. To estimate the EDR ($\epsilon^{1/3}$) in this inertial region, there are four basic steps, listed below.

1. Ensure the flight is within nominal characteristics.
2. Sample the wind speed and compute the energy spectrum.
3. Compare $E(k, \epsilon) = C\epsilon^{2/3}k^{-5/3}$ to a standard model, $E_0(k; 1)$, with $\epsilon = 1$.
4. Report EDR = $\epsilon^{1/3}$.

It is the sampling rate of the EDR that prevents the inclusion of structures that are either too large or too small. For the large scale, one assumes that the energy spectrum is a function of ϵ and E_T (rather than ϵ and k) and the balance of these processes implies characteristic scales of

$$\mathsf{L}_\infty \sim \epsilon^{-1}E_T^{3/2}, \qquad \mathsf{T}_\infty \sim \epsilon^{-1}E_T, \qquad \mathsf{U}_\infty \sim E_T^{1/2}, \qquad \mathrm{Re}_\infty = \epsilon^{-1}\nu^{-1}E_T^2. \tag{5.5}$$

As energy moves through the spectrum from large to small scales, it is eventually dissipated into heat through the viscosity. The size of the smallest eddies can be estimated by assuming that this is only determined by the dynamic viscosity ($[\nu] = \mathsf{m}^2\ \mathsf{s}^{-1}$) and ϵ. Accordingly, we have

$$\mathsf{L}_0 \sim \epsilon^{-1/4}\nu^{3/4}, \qquad \mathsf{T}_0 \sim \epsilon^{-1/2}\nu^{1/2}, \qquad \mathsf{U}_0 \sim \epsilon^{1/4}\nu^{1/4}, \qquad \mathrm{Re}_0 = 1. \tag{5.6}$$

As an example, consider $\epsilon \sim 0.2\ \mathsf{m}^2\ \mathsf{s}^{-3}$ with $\nu = 1.48 \times 10^{-5}\ \mathsf{m}^2\ \mathsf{s}^{-1}$ and $E_T = 20\ \mathsf{m}^2\ \mathsf{s}^{-2}$. For these nominal values, the large scale is characterized by

$$\mathsf{L}_\infty \sim 500\ \mathsf{m}, \qquad \mathsf{T}_\infty = 100\ \mathsf{s}, \qquad \mathsf{U}_\infty = 4.5\ \mathsf{m}\ \mathsf{s}^{-1}, \qquad \mathrm{Re}_\infty \sim 10^8, \tag{5.7}$$

while the small scale is characterized by

$$\mathsf{L}_0 = 3.5 \times 10^{-4}\ \mathsf{m}, \qquad \mathsf{T}_0 = 8.6 \times 10^{-3}\ \mathsf{s}, \qquad \mathsf{U}_0 = 4.2 \times 10^{-2}\ \mathsf{m}\ \mathsf{s}^{-1}. \tag{5.8}$$

The viscosity decreases with temperature, thus making all these scales smaller at higher altitudes.

In the current version of the NCAR system, a sampling frequency of $f = 8$ Hz is used. For a heavy aircraft flying at 250 $\mathsf{m}\ \mathsf{s}^{-1}$ (Mach 0.73), the length scale per sample is $250/(8/2) = 62.5$ m, whereas for a light aircraft flying at 75 $\mathsf{m}\ \mathsf{s}^{-1}$ (270 $\mathsf{km}\ \mathsf{hr}^{-1}$), this drops to 18.75 m. The key point is that every possible aircraft speed corresponds to a distance travelled per sample, L, that satisfies $10^{-3}\ \mathsf{m} < \mathsf{L} < 500\ \mathsf{m}$.

5.3 Research threads

In this section we detail all of the methods used during the workshop.

Statistical Model: The statistical properties of the EDR values stored in the IATA database were summarized. A seasonal pattern was detected using a linear GAM model.

Voronoi: To partition the domain into regions associated with each of the sampled EDR values, a Voronoi diagram was constructed. This does not seem to be viable because of the dense clustering near the airline hubs.

KDE Heat Map: Rather than partitioning the domain, another method is to diffuse effectively each EDR datapoint into its surrounding locality. The resulting extrapolated EDR field, obtained by employing a kernel density estimation, achieved better results than the Voronoi partition. The drawback was that regions where EDR datapoints overlapped still caused apparent overestimation of the EDR when considering large spatial scales.

Clustering: To remedy the issue of overlapping regions, we clustered points that are spatially close and temporally close. Using the cluster statistics together with the original points shows some promise of producing a useful forecast of EDR.

Spatio-temporal Weighting: We estimated EDR using the intuition that clustered points should be more heavily weighted than isolated points. This is a specific case of the clustering concept. By computing an overall background, which depends on the size of the domain, the overlapping problem with the kernel estimation was avoided. Setting a background cutoff level naturally excludes datapoints that are either old (denoted as stale) or regions that are too far away from any EDR datapoint. An animation of this implementation has been provided.

5.3.1 Summary statistics

The data used in this study was obtained from IATA and consists of $16,000,000$ observations over a time period from September 2020 to August 2022. Each year there are anywhere from $500,000$ to $800,000$ flights contained in the database.

To obtain a seasonal heat map, the $\{\mathrm{EDR}_i\}_{i=1}^N$ values are extracted with three associated predictor variables $\{(x_{1,i}, x_{2,i}, t_i, a_i)\}_{i=1}^N$, where x_1, x_2 correspond to the longitude and latitude of a geographic point, t denotes the time, and a denotes the altitude. A temperature sample was also available but was omitted because of the strong linear correlation between temperature and altitude. It was found that the response variable (EDR) was both zero-inflated and heavily skewed. All these data characteristics are shown in Figure 5.3.

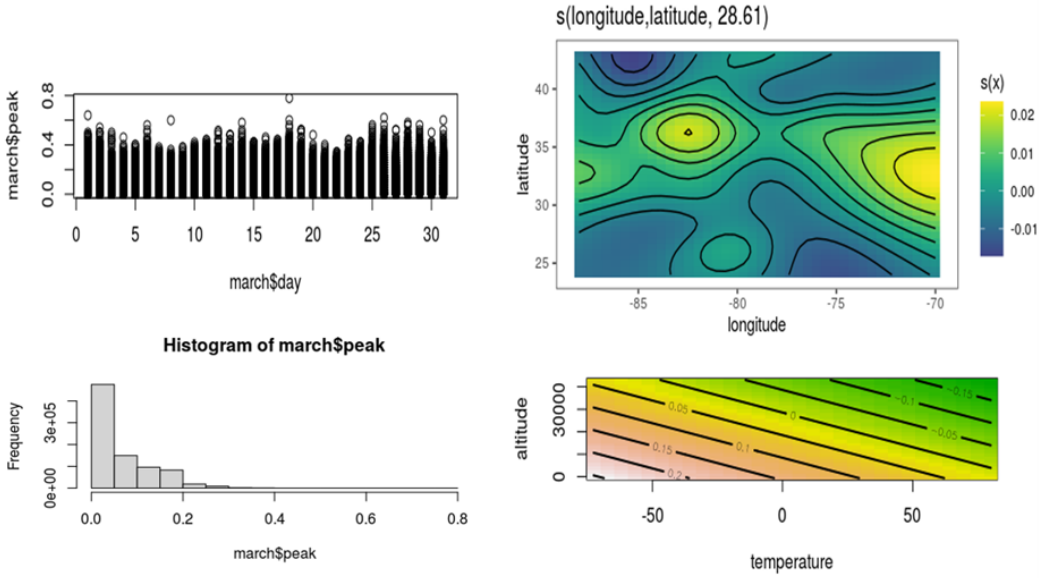
s(longitude,latitude, 28.61)

Histogram of march$peak

Figure 5.3: Overall statistical structure of the EDR samples in the IATA database.

5.3.2 Seasonal variation model

To explore the monthly effect on EDR, we used a generalized additive model (GAM), in which the response variable (EDR) depends linearly on unknown smooth functions of the positional variables. In particular, for each sample $(x_{1,i}, x_{2,i}, t_i, a_i)$, we have

$$\log(\mu_i + 1) = \beta_0 + \beta_1 t_i + \beta_2 a_i + s(x_{1,i}, x_{2,i}), \qquad y_i \sim \mathcal{N}(\mu_i + 1, \sigma^2), \tag{5.9}$$

where $\{\beta_j\}_{j=0}^2$ are fixed coefficients, $s(x, y)$ is a smooth function of the position (x, y), and μ_i is the expected EDR and $y_i = \mathrm{EDR}_i + 1$. A log-link function of the EDR is used, amounting to a power relationship. Figure 5.4 shows the seasonal variation of the EDR events over a single year.

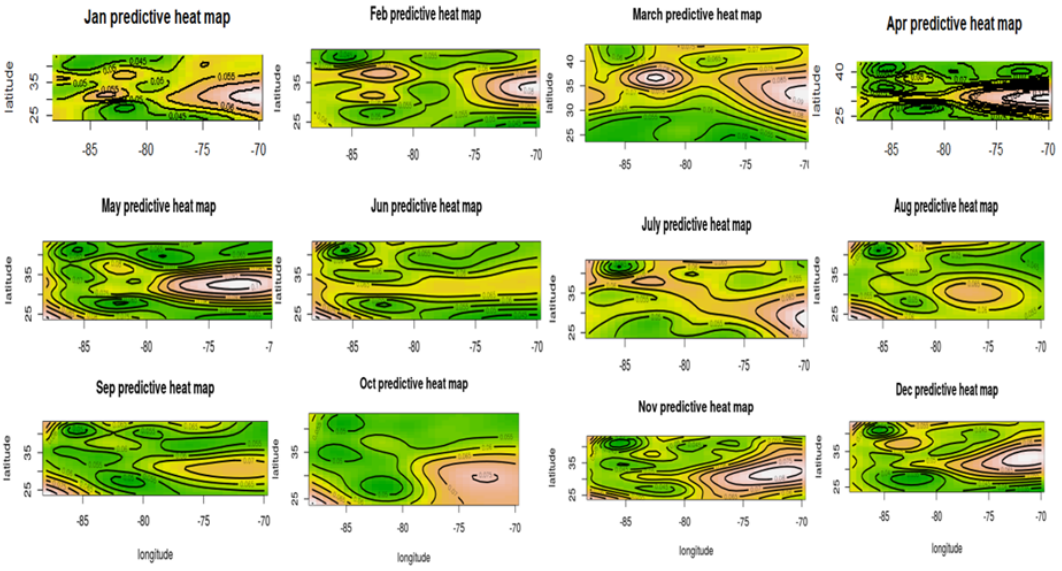

Figure 5.4: The seasonal heat maps for the EDR samples in the IATA database.

5.3.3 Seasonal variation results

Two main trends were detected within the seasonal data. (i) Altitude has a negative effect on EDR across all the months, i.e., increasing the altitude decreases the EDR. (ii) The data is more complex with EDR decreasing over the months May–September and November and increasing over December–April and October.

A 5-fold cross-validation was used for model comparison and to detect any over-fitting. This effect is likely to occur in models with complex structures. Our observations are that the average residual in each month is around 0.086. The error rate is close to 10%, which is relatively large. A possible reason for this is that the proposed distribution is incorrect. One should try to use other distributions for the response variable, such as the distributions in quasi-families. Another possible source of this discrepancy is the violation of the assumption of independent and identically distributed (iid) residuals. The spatial and temporal correlations between the residuals point to that explanation. The binning of the data into monthly blocks (referred to as faceted) can also cause discrepancies. For example, in March, we can see adjacent dates have similar residual distributions: in some areas there are always positive residuals and in other areas only negative residuals are observed. To get around this difficulty we can use a more sophisticated spatio-temporal model that will introduce prescribed correlations in space and time.

5.3.4 Voronoi diagram

To begin a Voronoi mapping was constructed utilizing EDR incidents in order to visualize turbulence across a region or "cell" based on Voronoi theory. The results of a simple Voronoi plot overlaid onto a world map with points can be seen in Figure 5.5.

Since reports of EDR measurements can be quite densely packed due to reporting standards (i.e., increased reporting frequency for higher EDR rates), the diagram did not end up helping to visualize the data. Furthermore the sizes of the Voronoi cells are based on the distances between EDR incidences and do not capture the real dynamics/patterns of turbulence. Though a heatmap could be overlaid, this would not help alleviate the core issues with Voronoi mapping.

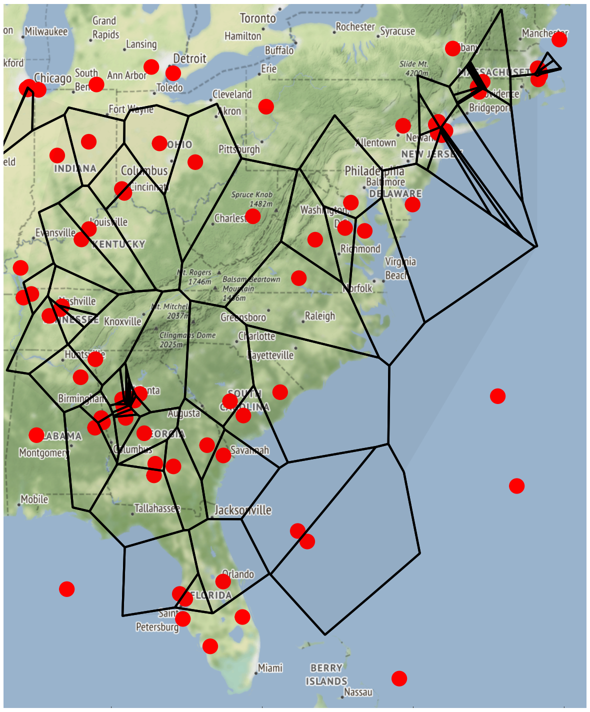

Figure 5.5: An example Voronoi diagram for a sample of EDR data.

5.3.5 Kernel Density Heat Map

Each of the EDR data points specify the turbulence at a specific point along the flight path. One of the fundamental problems is to infer the values of EDR in a neighbourhood of this position. A systematic way to accomplish this is to start with the cumulative distribution of the indicator function

$$F_X(x) = \mathbb{E}\left(\mathbf{1}_{\{X\leq x\}}\right), \qquad \mathbf{1}_{\{X\leq x\}} = \begin{cases} 1, & X \leq x, \\ 0, & \text{otherwise}, \end{cases} \tag{5.10}$$

which is naturally related to a histogram of the associated probability density function of the random variable X. The function $F_X(x)$ contains all of the information about X but it is not unique. An alternative way to describe X is through the characteristic function $\varphi_X : \mathbb{R} \to \mathbb{C}$ defined as

$$\varphi_X(t) = \mathbb{E}\left(\mathrm{e}^{itX}\right) = \int_{\mathbb{R}} \mathrm{e}^{itx}\, \mathrm{d}F_X(x), \tag{5.11}$$

showing the connection between φ_X and F_X.

Consider a collection of N points $\{x_1, x_2, \ldots, x_N\}$ taken from X that are all equally likely with a probability of $1/N$. In this case an estimate for the characteristic function is $\tilde{\varphi}(t) = (1/N)\sum_{j=1}^{N} \mathrm{e}^{itx_j}$. To obtain the probability density of $f_X(x)\,\mathrm{d}x = \mathrm{d}F_X(x)$ we can compute the Fourier transform of $\varphi_X(t)$ but with a nonnegative bandwidth limiting function $\lambda_h(t) = \lambda(ht)$. The role of $\lambda(t)$ is spread over the domain of each sample point. Associated with this are the conditions $\lambda(0) = 1$ and $\lambda(t) \to 0$ as $|t| \to \infty$. An estimate for $f_X(x)$ is then given by

$$\tilde{f}_X(x) = \frac{1}{2\pi}\int_{\mathbb{R}} \mathrm{e}^{-itx}\tilde{\varphi}(t)\lambda_h(t)\,\mathrm{d}t = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{2\pi}\int_{\mathbb{R}} \mathrm{e}^{it(x_j-x)}\lambda(ht)\,\mathrm{d}t = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{h}K\left(\frac{x_j-x}{h}\right), \tag{5.12}$$

where $K(x)$ is the Fourier transform of $\lambda(t)$. Popular choices for $\lambda(t)$ are

$$\lambda_1(t) = \begin{cases} 1, & |t| \leq 1, \\ 0, & \text{otherwise}, \end{cases} \quad K_1(x) = \frac{\sin x}{\pi x}, \quad \lambda_2(t) = \mathrm{e}^{-\alpha t^2}, \quad K_2(x) = \frac{1}{2\pi}\left(\frac{\pi}{\alpha}\right)^{1/2}\mathrm{e}^{-x^2/4\alpha}. \tag{5.13}$$

The Kernel Density Heat Map generalizes this process to the two-dimensional spatial coordinates with a covariance matrix C playing the role of h, so that

$$\tilde{f}_X(\mathbf{x}) = \frac{1}{N}\sum_{j=1}^{N} |\mathsf{C}|^{-1/2} K\left(|\mathsf{C}|^{-1/2}(\mathbf{x}_j - \mathbf{x})\right). \tag{5.14}$$

In Figure 5.6 this is implemented in R using `R-Shiny` and the smoothing function is a Gaussian in each coordinate.

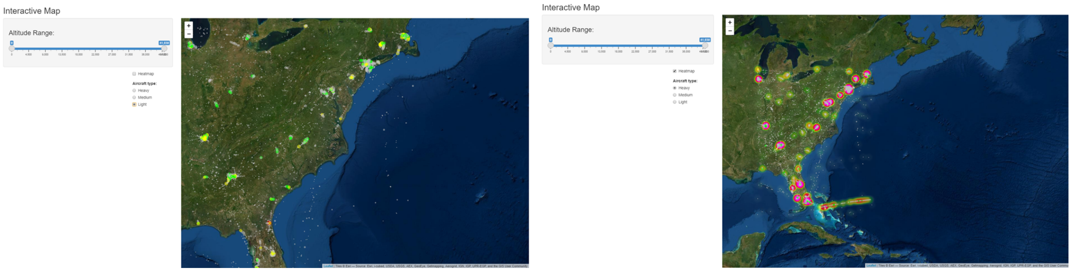

Figure 5.6: A simple heat map implementation for the EDR based on location and using a Gaussian smoothing function in each coordinate.

5.3.6 Heat map clustering

In Figure 5.7 we illustrate a simple clustering approach. EDR points are separated into “zones” based on altitude strips of 4,000 ft and a time frame of four hours, shifting throughout the range for each of these variables. After separating the data into these different zones, clusters were created based on the distance between EDR points. Those that were within one degree of longitude and latitude were included into the respective clusters. After partitioning all points within a zone into clusters, the cluster statistics were recorded for use *in conjunction* with the original points. For Figure 5.7 the statistic that was kept was simply the centre location of all the points in a given cluster, with a circle radius based on EDR intensity. This figure illustrates a concept since any parameter configuration based on cluster statistics could be used, taking into account a variety of aspects such as positioning, shape, or size.

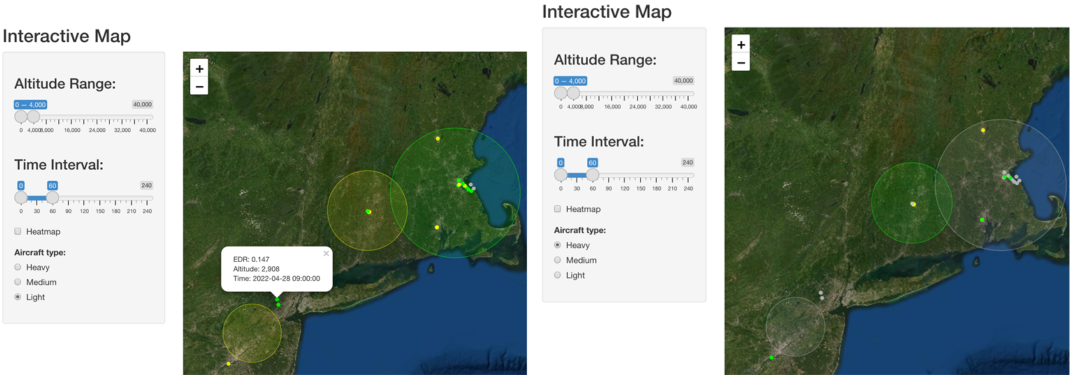

☐ Heatmap

Aircraft type:
○ Heavy
○ Medium
● Light

EDR: 0.147
Altitude: 2,908
Time: 2022-04-28 09:00:00

☐ Heatmap

Aircraft type:
● Heavy
○ Medium
○ Light

Figure 5.7: Results of the clustering of coordinates using a circle radius which is a function of the EDR.

5.3.7 Spatio-temporal weighting

In this section we consider an approach that will interpolate a set of datapoints with the additional aspect of only including data that was sampled within a particular time interval. Consider a dataset $D = \{(\bar{\mathbf{x}}_i, \bar{t}_i, \epsilon_i)\}_{i=1}^{r}$ of position-time-EDR measurements over a time interval $[0, T]$. For each datapoint assign the weighting function

$$w_i(x,t) = \mathrm{e}^{-\|\mathbf{x}-\bar{\mathbf{x}}_i\|/\xi}\mathrm{e}^{-(t-\bar{t}_i)/\tau}\mathbf{1}_{\{\bar{t}_i \le t\}}(t), \qquad i = 1,2,\dots,r, \tag{5.15}$$

where τ and ξ are characteristic time and length scales, respectively. The indicator function $\mathbf{1}_{\{\bar{t}_i \le t\}}(t)$ naturally excludes any data subsequent to time t, while the value of ξ (resp. τ) defines a displacement (resp. interval) where the weight drops by a factor of e^{-1}. With this weighting function the EDR field is defined only at those points $(\mathbf{x},t)$ where the total weight exceeds some critical value w^*. This means that the EDR field is defined as $\tilde{\epsilon} : \Omega \to \mathbb{R}^+$ with

$$\tilde{\epsilon}(\mathbf{x},t) = \frac{\sum_{i=1}^{r} w_i(\mathbf{x},t)\epsilon_i}{\sum_{i=1}^{r} w_i(\mathbf{x},t)}, \qquad \Omega = \left\{ (\mathbf{x},t) \,|\, \sum_{i=1}^{r} w_i(x,t) > rw^* \right\}. \tag{5.16}$$

The parameter τ together with w^* dictates a time interval beyond which the datapoints are considered stale and eliminated from the EDR field. In a similar way, ξ and w^* dictate the maximum allowable spatial distance from any given datapoint.

Computational implementation

The weighted EDR $\tilde{\epsilon}$ is now a continuous scalar field defined over a subset of the geographical region $[\mathrm{lat}_{\min}, \mathrm{lat}_{\max}] \times [\mathrm{long}_{\min}, \mathrm{long}_{\max}]$ and a time interval $[0,T]$. Therefore a two-dimensional heat map can be constructed by plotting a large amount of level sets.

Consider a two-dimensional grid of Ω made of N-by-M cells (spatial mesh). The mesh, therefore, contains $(N+1)(M+1)$ points, which can be arranged in a vector $\mathbf{x} = \{x_j\}_{j=0}^{(N+1)(M+1)-1}$. In addition we consider discrete time steps $\{t_n\}_{n=0}^{L}$, where $t_0 = 0$, $t_L = T$, and in general $t_n = n\Delta t$. The time resolution is therefore $\Delta t = T/L$. First we compute the sum in the denominator of (5.16), where for each (x_j, t_n) we have

$$s(x_j,t_n) = \sum_{i=1}^{r} \mathrm{e}^{-(t_n-\bar{t}_i)/\tau}\mathbf{1}_{\{\bar{t}_i \le t_n\}}(t_n)\mathrm{e}^{-\|x_j-\bar{x}_i\|/\xi}, \tag{5.17}$$

or in matrix form with $[S]_{jn} = s(x_j,t_n)$, $A \in \mathbb{R}^{(L+1)\times r}$, $B \in \mathbb{R}^{r\times(N+1)(M+1)}$,

$$S = AB, \qquad [A]_{ni} = \mathrm{e}^{-(t_n-\bar{t}_i)/\tau}\mathbf{1}_{\{\bar{t}_i \le t_n\}}(t_n), \qquad [B]_{ij} = \mathrm{e}^{-\|x_j-\bar{x}_i\|/\xi}. \tag{5.18}$$

For the numerator of (5.16) we have

$$s_\epsilon(x_j,t_n) = \sum_{i=1}^{r} \mathrm{e}^{-(t_n-\bar{t}_i)/\tau}\mathbf{1}_{\{\bar{t}_i \le t_n\}}(t_n)\epsilon_i\mathrm{e}^{-\|x_j-\bar{x}_i\|/\xi}, \tag{5.19}$$

or if computed as the product of three matrices,

$$S_\epsilon = AEB, \qquad E = \mathrm{diag}(\epsilon_1, \epsilon_2, \dots, \epsilon_r) \in \mathbb{R}^{r\times r}. \tag{5.20}$$

Therefore each entry of the discrete version of the weighted EDR function $\tilde{\epsilon}$ can be computed as the quotient of the corresponding entries of S_ϵ and S whenever $S_{jn} > rw^*$.

Testing

To test this approach with different sets of data, we consider a 48-hour interval of datapoints (a typical dataset would only contain 4 hours). The parameters used to compute the EDR field are given in the table. The results corresponding to these parameters are available in an .mp4 file that accompanies this report. Some snapshots are shown in Figure 5.8.

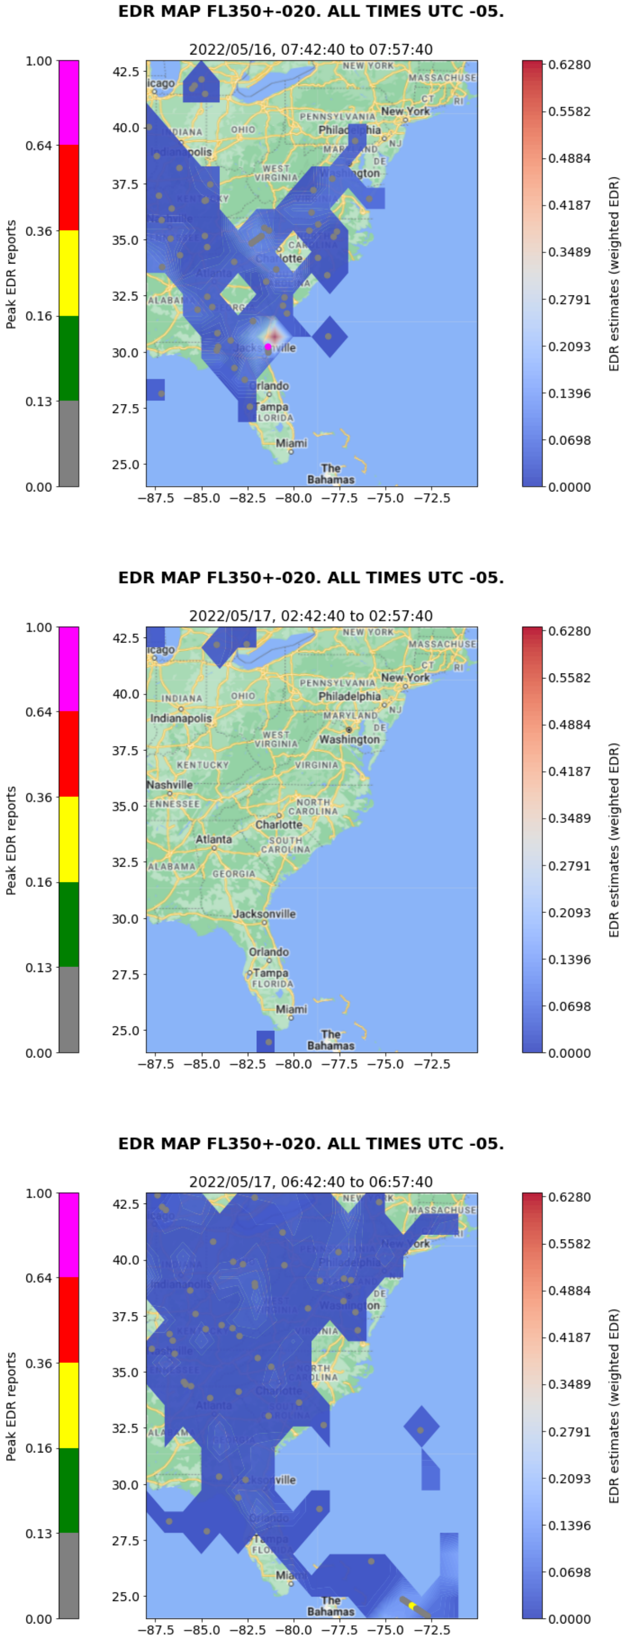

Figure 5.8: Snapshots of the weighted EDR at different times showing the ability to predict high EDR values in regions with high EDR measurements (top), the ability of not predicting EDR values where there are no datapoints (middle), and a typical situation where there is a mix of low and medium intensity EDR values (bottom).

Data parameters	
Start time	May 16, 2022, 07:42:40 (UTC -05)
End time	May 18, 2022, 07:42:40 (UTC -05)
Range of latitudes	min : 24.0179°, max : 43.0°
Range of longitudes	min : −87.993°, max : −70.05°
Altitude	35,000 ft ± 2,000 ft
Model parameters	
Characteristic length	$\xi = 0.065^\circ$
Characteristic time	$\tau = 10$ minutes
Threshold	$w^* = 10^{-12}$
Spatial resolution	$N = 18, M = 20$
Time resolution	$\Delta t = 15$ minutes
End time	$T = 48$ hours

5.4 Conclusions and future work

A multi-pronged approach to the building of a forecasting map of turbulence yielded a number of conclusions. These are itemized below for the various sections of the report.

Statistical Approach

- Statistics are based on 16 million observations over a two-year period.
- A strong correlation between altitude and temperature validates the data.
- The probability distribution of EDR values is strongly skewed.
- Seasonal variations were detected both in EDR intensity and where large intensities are more likely.

Voronoi Partitions

- The Voronoi method of partitioning the domain suffers from the dense clustering of the points.
- Datapoints along a specific trajectory naturally lend themselves to bandwidth limited methods, rather than domain decomposition.

Interactive KDE Heat Map

- allows dynamic changes of the altitude slice and the time interval
- provides a systematic way to choose kernels for both the spatial and temporal locations

Spatio-temporal Weighted Heat Map

- The method lends itself to a rapid solution through matrix multiplication.
- Regions of low confidence are efficiently excluded with a cutoff condition.
- In the implementation the spatial and temporal weights are exponential decays but can be easily changed.

After spending a week examining the IATA database and its collection of EDR observations, a number of issues remain outstanding and should be addressed. With such a large database of instances the statistical analysis can rely on the dataset itself, rather than using the assumption of normality. It is expected that this will enhance the ability to observe the seasonal variations. The seasonal variations will not be a significant effect for any given flight. Currently the distances are assumed to be uniform in latitude and longitude. In reality, for a fixed longitude $\theta = \theta_0$, the differential length along the longitude is $\Delta s_{\text{long}} = R\Delta\phi$, where $\Delta\phi$ is a change in latitude. For a fixed latitude $\phi = \phi_0$, the differential length along the latitude is $\Delta s_{\text{lat}} = R\cos\phi_0\Delta\theta$, where $\Delta\theta$ is a change in longitude. Note that we have $\Delta s_{\text{lat}} = \Delta s_{\text{long}} \cos\phi_0$ and only along the equator where $\phi_0 = 0$ are these two distances equal.[4] These considerations will not affect the algorithms used to generate an EDR map but will distort any two-dimensional projection.

The chosen implementation of a spatio-temporal model to calculate the EDR map uses ad hoc exponential weighting functions. These should be chosen in a way that is consistent with the required band limiting. The band limiting heuristic should also be determined by the local region by performing an initial clustering based on a variety of characteristics. For example, EDR level, cluster radius, and geographic position could be used. Indeed the reported EDR $= \epsilon^{1/3}$ is proportional to the standard deviation of the eddy size in any particular turbulent field. With this initial clustering pass, the parameters ξ, τ, and w^*, could be tuned dynamically with the arriving data. One thing that was not considered was the evolution of the EDR field between individual samples. While this is an important aspect, an accurate solution of a turbulence model (for example k-ϵ) for the air flow requires an assumption on the initial state of the EDR field. This is what we attempting to model, however, and the fact that we were looking at four-hour intervals more or less guarantees that the unknown field will not appreciably change over an interval.

Executive summary

The 12th Montreal Industrial Problem Solving Workshop held at the Université de Montréal, on August 22–26, 2022, concerned itself (among other problems) with the IATA question of the possibility of "Creating a Heat Map and Building a Seasonal Diagram." We feel that, based on the observations and analysis of this report, the EDR data has high value and within them is the potential to create a sophisticated forecast of clear air turbulence.

For the data that was supplied, a seasonal variation in the intensity of the eddy dissipation rate (EDR) was detected. With this variation, there is also a predictable geographic variation of the pattern. This second observation is more subtle and it is expected that it could be modelled if the statistics defined by the observations were used instead of assuming normality. Another feature of the current statistical model is that EDR values are not uniformly sampled. Regions of large EDR are avoided when detected and this will implicitly bias the database.

In parallel with the statistical aspects, a majority of the effort was concerned with the viability of producing a map of the turbulence. The current in-situ method estimates the EDR by sampling the wind velocity and computing a spectrum of the energy. The sampling frequency is chosen so that both heavy and light aircraft are guaranteed to be sampling the spectrum from within the inertial subrange. In addition, to minimize the cost of transferring the data to the ground, EDR values are not reported unless they exceed a certain threshold for a particular duration. Bounds on the unreported EDR values can be derived but this is nontrivial given that the current error suppression algorithm is nonlinear. More time would be required to do this: perhaps also required are updates of the current error encoding algorithm to simplify inversion.

One of the fundamental issues with the datapoints is that they are located along the flight path rather than being distributed randomly throughout the domain. One way to determine an EDR map would be to partition the domain into those regions nearest to the given datapoints. This is not viable with the EDR samples, specifically with the dense clusters of points near the airline hubs, and another method was required. An alternative to partitioning the domain is to diffuse the datapoints into the neighbouring locations. This is done by using an appropriate kernel that spreads the data. A few examples have been given but there is a lot of flexibility with this choice. A naive choice of a Gaussian kernel in space has the artifact that when large domains are chosen, as when the map is zoomed out, the average EDR seems to increase. It does this because the average EDR value is assumed to be spread over the whole domain as it enlarges.

---

[4] One degree of latitude equals $\frac{\pi}{180}R$ and thus approximately 364,000 feet (69 miles). One degree of longitude equals $\frac{\pi}{180}R\cos\phi$ and thus depends on the latitude, ϕ, and the radius of the earth, R.

The issue with overlapping regions can be corrected if the sample is normalized. Such an implementation is given in the animation attached to the current report. In this case the datapoints are assumed to decay exponentially in space and time. This last effect is necessary to simulate a systematic lack of confidence in datapoints as they age. The provided algorithm can be easily implemented on a rudimentary system as only matrix multiplication and lookup tables are required to compute the map.

More sophisticated models are possible with an initial clustering of the data with respect to a variety of features. Carrying this out as a preprocessing step would create various “zones” for a map. Currently a zone is a region around each datapoint, independent of any other features. Within each zone, a map could be computed based on the local features: the “zone maps” could then be patched together to build a final map.

Bibliography

[1] Federal Aviation Administration `https://www.faa.gov/travelers/fly_safe/turbulence`.

[2] Meymaris G., Sharman R., Cornman L.B., Deierling, W., (2019). The NCAR In Situ Turbulence Detection Algorithm, NCAR/TN-560+EDD, 1–53.

[3] Sharman R.D., Cornman L.B., Meymaris G., Pearson J., Farrar T., (2014). Description and Derived Climatologies of Automated In Situ Eddy-dissipation-rate Reports of Atmospheric Turbulence, Journal of Applied Meteorology and Climatology, 53(6): 1416–1432.

6 Automatic text simplification of news articles in the context of public broadcasting

Diego Maupomé [a]
Fanny Rancourt [a]
Thomas Soulas [a]
Alexandre Lachance [a]
Guillaume Poirier [a]
Marie-Jean Meurs [a]
Desislava Aleksandrova [b]
Olivier Brochu Dufour [b]
Igor Pontes [b]
Rémi Cardon [c]
Michel Simard [d]
Sowmya Vajjala [d]

[a] *Université du Québec à Montréal*
[b] *CBC/Radio-Canada*
[c] *Université catholique de Louvain*
[d] *Conseil National de Recherche du Canada*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Maupomé, Rancourt, Soulas, Lachance, Meurs, Aleksandrova, Brochu Dufour, Pontes, Cardon, Simard, Vajjala

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

6.1 Introduction

This report summarizes the work carried out by the authors during the *Twelfth Montreal Industrial Problem Solving Workshop*, held at the Université de Montréal in August 2022. The team tackled a problem submitted by CBC/Radio-Canada on the theme of *Automatic Text Simplification* (ATS). In order to make its written content more widely accessible and to support its second-language teaching activities, CBC/RC has recently been exploring the potential of automatic methods to simplify texts. They have developed a modular *lexical simplification system* (LSS), which identifies complex words in French and English texts and replaces them with simpler, more common equivalents. Recently, however, the ATS research community has proposed a number of approaches that rely on deep learning methods to perform more elaborate transformations, not limited to lexical substitutions, but including syntactic restructuring and conceptual simplifications as well. The main goal of CBC/RC's participation in the workshop was to examine these new methods and to compare their performance with that of their own LSS.

This report is structured as follows. In Section 6.2, we detail the context of the proposed problem and the requirements of the sponsor. We then give an overview of current ATS methods in Section 6.3. Section 6.4 provides information about the relevant datasets available, both for training and testing ATS methods. As is often the case in natural language processing applications, there is much less data available to support ATS in French than in English; therefore we also discuss in that section the possibility of automatically translating English resources into French, as a means of supplementing the French data. The outcome of text simplification, whether automatic or not, is notoriously difficult to evaluate objectively; in Section 6.5, we discuss the various evaluation methods we have considered, both manual and automatic. Finally we present the ATS methods we have tested and the outcome of their evaluation in Section 6.6. Section 6.7 concludes this report and presents research directions.

6.2 Problem description

6.2.1 Context

The mandate of CBC/Radio-Canada is to inform, enlighten, and entertain all Canadians. When presenting its plan on equity, diversity, and inclusion for 2022-2025, CBC/Radio-Canada committed itself to doing the utmost so that all persons living in Canada feel valued, recognized, and heard by their public broadcaster from sea to sea. On its website (radio-canada.ca), Radio-Canada publishes between 450 and 600 articles each day. These articles deal with complex topics (health crisis, climate crisis, the economy, the polarization of society, international conflicts, etc.). Since a good understanding of current issues is necessary to take part in the democratic debate, Radio-Canada thinks that the use of ATS could help a greater number of citizens take part in this debate. Simplifying or summarizing some of its written contents (in an automatic fashion) could enhance the understanding of the articles and make them more attractive for people struggling with literacy, neurodiverse people, and new immigrants (for instance).

In April 2021 CBC/Radio-Canada launched Mauril[1], a digital platform for learning French and English through audio-visual information produced by the public broadcaster. The development team is currently trying to broaden the supply of written content through a new reading comprehension task. To this end Radio-Canada will need ATS to produce simpler versions of original articles, in order to take the level of beginners (learning French or English) into account.

---

[1] `https://mauril.ca/fr/`

6.2.2 The problem

ATS consists in decreasing the complexity of a text, from the lexical and syntactic points of view, while retaining its meaning and grammaticality. This has been shown to improve the readability and ease of understanding of the text. Methods for simplifying text fall into two categories: modular systems (which carry out lexical or syntactic simplifying operations iteratively or recursively), and end-to-end systems, which learn to carry out several modifications at the same time through labeled data. In modular systems the transformations are in most cases applied sentence-wise.

6.2.3 Desired solution

We wish to compare the performance of our modular lexical simplification systems with that of an end-to-end system that operates at the sentence level, without being limited to lexical simplifications. The supervised end-to-end approaches that we have identified make it possible to perform several types of transformations (both syntactic and lexical) but require a large number of examples. On the other hand at least two unsupervised approaches [13, 25] make it possible to learn how to simplify from free text, without depending on aligned corpora. We also need to identify the appropriate evaluation metrics and evaluation corpora in both French and English.

6.3 Related work

Automatic textual simplification consists in reducing the complexity of a text (in terms of lexicon and syntax) while retaining its original meaning, in order to improve its readability and understanding. In reality, textual simplification is most often achieved by transformations performed at the sentence level. These are rewriting transformations such as replacement (lexical simplification), reorganization (syntactic simplification), and splitting. In accordance with the literature, our investigation was centered around sentence-level simplification.

Broadly construed, ATS systems will seek to parse an input sentence and produce an output sentence presumed to be equivalent to it but linguistically simpler. As previously mentioned, the inner workings of these systems place them into one of two broad categories: modular and end-to-end systems. Modular systems will apply a set of definite, linguistically informed transformations. These transformations may pertain to different aspects of language and simplification. For example one transformation may seek to replace words deemed complex with simpler ones, while others will operate on the syntactic structure of sentences. Assuming the independence of these aspects as they pertain to simplicity, these transformations will be applied independently by such systems.

In contrast end-to-end systems will seek to apply such transformations jointly. Importantly, said transformations are not necessarily explicitly parameterized by these approaches. Instead, end-to-end models will transform the input sentence into a dense vector representation from which the output sentence will be produced, and are thus similar to encoder-decoder models. These neural network models will be trained to match the expected output; the simplifying transformations are expected to be learned and effected implicitly through the inner representations of the network.

Thus end-to-end approaches will usually require training on parallel corpora, i.e., datasets matching complex sentences with their simplified versions. While it is conceivable to train simple so-called sequence-to-sequence models on this task, their high capacity makes them prone to overfitting on the smaller datasets habitually available in ATS. As such, end-to-end approaches based on inference from parallel corpora will attempt to embed structural constraints in the encoding-decoding process in order to rein in model capacity. For example, the sentence simplification task can be reframed by having the model issue edit operations to be applied on the input sequence rather than produce the output sequence outright. These operations, once interpreted and realized, would result in the simplified sentence. For example, the model might be tasked with issuing deletion and preservation operations for each token in the input sentence. Carrying out these instructions on every token would result in a simplified sentence. The true edit operations constituting the expected output of the model can be produced automatically from the aligned sentence pairs. The output space of this modified framework is greatly reduced in dimension, thereby simplifying the task. This approach has been carried out in the literature, achieving competitive results [8, 15].

Rather than mitigating the impact of limited training examples by reframing the model task, one can train encoder-decoder models directly to produce simplified text on non-parallel corpora. These corpora comprise sets of sentences of each complexity level that are not individually paired. Thus large amounts of data are more easily gathered. Nonetheless such data do not give a mapping from a complex to a simplified sentence and the model cannot be trained in a supervised manner. Surya et al. [25] proposed to circumvent this issue by deploying an encoder-decoder model with two decoders, one for each of two complexity levels. These decoders are subject to separate training criteria. In particular the decoder producing simple outputs must process both complex sentences and simple sentences in similar manners while doing so differently from the complex decoder. Together the encoder and the simple decoder constitute a simplification model. With this approach one obtains competitive results on standard datasets. The authors, however, note that performance can be further improved by additional training on parallel data.

Another approach to addressing limited data is to pose text simplification as a particular case of paraphrasing [13]. Paired paraphrase examples can be extracted automatically in larger numbers as there is no need for any assessment of their difference in terms of complexity. Of course, in order to use this model for simplification, one must be able to tilt the paraphrasing model towards the production of a simpler paraphrase. To this end, the paraphrase model is trained with auxiliary control inputs parameterizing the difference between the input and output sentence with regard to selected aspects [12], e.g., change in total length, change in dependency-tree depth. Once the model is deemed proficient at paraphrasing, simplifying values of these control inputs are determined empirically from parallel simplification corpora. This approach achieves a state-of-the-art performance on several datasets as well as good human ratings in English, French, and Spanish.

6.4 Datasets

ATS corpora vary in several respects, as each attempts to address specific preoccupations of ATS research. The modular lexical simplification system developed in this work relies on lexical resources exclusively for evaluation. In contrast the end-to-end sentence simplification approaches studied make use of sentence-level corpora both for training and for evaluation. Most resources available of either type are for English; very few resources exist for French.

6.4.1 Lexical simplification

A multilingual lexical simplification dataset was produced for the TSAR-2022 Shared Task in English, Spanish, and Portuguese [22]. The gold test set in English contains 373 sentences with an identified complex word and multiple simplification suggestions provided by annotators (25 or 26 in some cases). To our knowledge the only evaluation dataset for lexical simplification in French is presented in [9].

6.4.2 Sentence simplification

The resources of greatest interest in the development of sentence-level, end-to-end systems are parallel corpora. Broadly speaking, these datasets match equivalent text fragments of different levels of complexity. Nonetheless they may still vary in several respects. A primary distinction is their granularity, for example, whether they focus on sentence-level or larger-level simplification. Also datasets may differ in their characterization of the (relative) complexity of the units (e.g., sentences, paragraphs) within each pair thereof. For example, in some datasets, observations might consist of two sentences, one deemed complex and the other simple. Alternatively, they may belong to precisely defined categories, borrowed, for example, from Second-Language Acquisition standards. This is of particular interest in developing simplification systems capable of producing tailored simplifications. Additionally corpora may have several references per example per target level, offering different, presumably equivalent simplifications of one single sentence. This is particularly important in the evaluation of systems, casting a wider net over possible system outputs. Finally some corpora will characterize the differences between units in an observation with respect to several aspects, such as how well the (semantic) meaning is preserved throughout or whether units are grammatically sound.

6.4.3 Training data

As mentioned, very few resources exist for text simplification in French. The only available resources were therefore reserved for evaluation. End-to-end approaches relying on parallel corpora were trained on a machine-translated version of the WikiLarge corpus [29]. This dataset is built from sentences extracted from the open collaboration online encyclopedia, Wikipedia, which has both English[2] and simplified English[3] (termed Simple English) versions. Sentences were extracted and aligned automatically from articles existing in both versions, for a total of 296k pairs. WikiLarge is the largest parallel ATS corpus, to our knowledge. Its translation was carried out using the Helsinki machine translation model [16]. More details about this process and its considerations are presented in Section 6.4.5.

In contrast the UNTS approach [25] does not require a parallel corpus. Rather sentences are taken from two separate sets of sentences of the desired levels of complexity, easing the requirements of data collection. Two sets of data were collected from different sources for the present work. Both are news sources, in order to mitigate the impact of domain shift from one set to another. Complex sentences were taken from the French-language portion of the MLSUM corpus [20], a multilingual dataset for automatic summarization. Simple sentences were collected from the simplified version of the Radio France Internationale (RFI) website[4] We term this collection RFI.

The datasets are listed in Table 6.1.

Table 6.1: Sentence-level resources used in training. WikiLarge was automatically translated into French.

Dataset	Language	Size	Description
WikiLarge	English	296K	Automatic alignment of English Wikipedia and Simple English Wikipedia
MLSUM	French	425K	Multilingual summaries of news
RFI	French	105K	Transcripts of a news feed in simple French

6.4.4 Evaluation data

Four datasets were considered for the evaluation of end-to-end sentence simplification systems: ALECTOR [10], OneStopEnglish [26], TurkCorpus [27], and ASSET [1]. The simplifications of all datasets are manual. Both ALECTOR and OneStopEnglish comprise document-level simplifications by experts, while TurkCorpus and ASSET comprise sentence-level simplifications by crowdsourced workers. The only French-language dataset among those selected is ALECTOR: the other datasets contain English-language text only. Our preliminary results concern only the ALECTOR corpus.

The **ALECTOR** corpus [10] comprises 79 excerpts aimed at children between the ages of 7 and 9, drawn from fiction and scientific texts common to the French curriculum. Each excerpt has one associated simplified version aimed at children with dyslexia or poor reading ability. Simplifying transformations, carried out by experts, included lexical simplification, deletion, and sentence splitting and merging. For our purposes, the dataset was adapted by manually aligning sentences. Complex

---

[2] `https://en.wikipedia.org/`

[3] `https://simple.wikipedia.org/`

[4] `https://francaisfacile.rfi.fr/`

excerpts in the original dataset were split at periods and colons, then manually aligned with their simplified counterparts, yielding 1165 pairs. Because the original simplification was carried out at the excerpt level, however, 23% of those pairs have no simplification.

The **OneStopEnglish** corpus [26] contains 2166 groups of sentences taken from OneStopEnglish lessons for English second-language learners. The sentences were originally sourced from news articles from The Guardian[5]. These articles were adapted for three different reading levels. From this text-aligned version sentence pairs were extracted from each level pair by automatically aligning sentences one-to-one. Pairs of sentences deemed too similar or dissimilar were discarded, resulting in different pair counts for each pair of levels.

The **TurkCorpus** [27] comprises 2359 English sentences from Wikipedia for which 8 independent simplifications were produced by different crowdsourced annotators. These annotators were instructed to rewrite sentences by substituting challenging words or idioms but without any content deletion or rearrangement, thus limiting the array of simplifying transformations present in the corpus.

To address these limitations, the **ASSET** corpus [1] provides 10 simplifications of the same source sentences as the TurkCorpus with a richer set of simplifying transformations: paraphrasing, lexical simplification, deletion, and sentence splitting and reordering.

6.4.5 Dataset translation

Due to the limited availability of French-language corpora, several datasets were automatically translated from English into French both for training and evaluation purposes. This was carried out using the Helsinki machine translation model [16]. Although the quality of machine translation has drastically improved over recent years, it is still important to survey the quality of the translation carried out on the corpora used. In particular it is important to study whether complexity was preserved through translation or, more precisely, whether the discrepancies in complexity within aligned pairs were preserved. Given that measuring linguistic complexity is a difficult problem in itself, however, certifying this preservation is difficult. Little work on these considerations can be found in the literature: namely Rauf et al. [19] sought to analyze the effects of translation upon the values of surface metrics at different complexity levels. Following this idea we measured changes within sentence pairs of the Wikilarge dataset using different metrics. As a measure of semantic proximity, we computed the cosine similarity between the LASER embeddings [5] of simple and complex sentences. Then the cosine similarity of each original pair was subtracted from that of the translated pair. The histograms are shown in Figure 6.1. Although most translated pairs have higher similarity than their original counterparts (56%), the difference in those cases appears to be small (median=0.012). The syntactic similarity, however, as approximated by the character-level edit distance, appears to diminish with translation: the median difference among translated pairs that are closer than the original (20%) is -4.

6.5 Evaluation

This section describes the evaluation methods used in the literature for the tasks of lexical and sentence simplification. The main components of text simplification are meaning preservation, grammaticality, and simplicity. Usually human evaluation explicitly uses these components while automatic metrics[6] can cover multiple aspects.

6.5.1 Human evaluation

Because it is difficult to model text simplification with mathematical quantities, human evaluation remains the gold standard. This is carried out along the aforementioned features with Likert scales.

---

[5] https://www.theguardian.com/

[6] In this report, the term "metric" is used in a broader sense than its mathematical definition. Thus evaluation scores that do not verify the triangular inequality are still referred to as metrics.

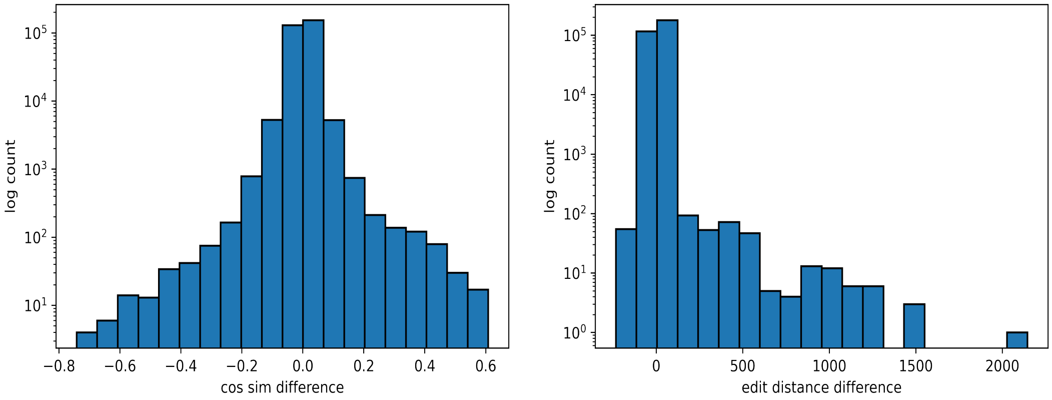

Figure 6.1: Left: Distribution of difference in cosine similarity between machine-translated (French) and original pairs in the Wikilarge dataset. Right: Distribution of difference in edit distances.

While the evaluation of grammaticality and meaning preservation is fairly straightforward, there are multiple approaches to the manual evaluation of simplicity: namely the simplicity of system outputs can be evaluated by themselves [23] or by comparing them with the input [11, 24, 27].

6.5.2 Automatic evaluation

Lexical simplification

The following evaluation metrics are used in the TSAR-2022 Shared Task for lexical simplification (assuming a correctly identified complex word is provided):

- Accuracy@1: the percentage of instances for which the best ranked substitution generated by the system is the same as the most frequently suggested simpler synonym in the gold data;
- Mean Average Precision@K: K$\in \{1,3,5,10\}$. MAP@K evaluates the relevance of the predicted substitutes and the position of the relevant candidates compared with the gold annotations;
- Potential@K: K$\in \{1,3,5,10\}$. Potential@K evaluates the percentage of instances for which at least one of the substitutions predicted is present in the set of gold annotations;
- Accuracy@K@top1: K$\in \{1,2,3\}$. ACC@K@top1 evaluates the ratio of instances where at least one of the K top predicted candidates matches the most frequently suggested substitute in the gold list of annotated candidates.

Sentence simplification

ATS has been frequently approached as a monolingual translation task [8, 29]: that is, the complex text is said to be translated into a simpler text within the same language. This formulation is manifest not only because of the parallels between approaches but also because work in the evaluation of ATS systems has been borrowing from the Machine Translation (MT) literature. Some metrics are borrowed without alteration, while other (ATS-specific) metrics draw a looser inspiration from MT metrics. Throughout the three key dimensions that these metrics seek to capture are:

- grammaticality: whether the output produced by some system is grammatically sound;
- meaning preservation: whether this output preserves the (semantic) meaning of the input; and
- simplicity: whether the output is indeed simple or simpler than the input.

Five metrics were selected for evaluation in the present work.

The **BLEU** metric [17], ubiquitous in MT, is equally common in ATS. It has notable weaknesses in this setting, such as its tendency to punish sentence splitting. It has, however, the advantage of also being usable in a reference-less setting by comparing to the input. Naturally this will reward less intervention. In this work, however, BLEU was used only to match against the expected output. Another MT metric, **BERTScore** [28], relies on pairwise similarities between the token representations of the system output and the references. These representations are computed from the BERT language model [7]. Basing the comparisons on a large language model rather than symbolic representations allows the metric to account for polysemy and makes it robust to paraphrasing. It has shown good correlation with human judgment in ATS tasks [21].

SARI [27], a metric specific to ATS, compares the output to both the reference and the input, in a manner similar to BLEU. Although very common in the ATS literature, it is difficult to interpret, and its grasp on grammaticality and meaning preservation remains unclear [21]. **SAMSA** [24], another ATS metric, seeks to measure the simplicity of outputs without the use of references by parsing for the semantic events described therein. It thus favours splitting sentences as well as reordering sentences into the dominant sentence structure for the language (Subject-Verb-Object for English). Lastly, the **ISiM** metric [14], also reference-less, computes the simplicity of a sentence by relying on the large-scale frequency of words as a proxy for their simplicity, i.e., texts with common words will be deemed simpler regardless of their meaning or lack thereof.

Except for ISiM, all metrics are supported by the EASSE package[7] [2], which was selected for use in this work. In future work, we may consider using a measure of similarity (e.g., cosine) between sentence embeddings to capture meaning preservation. Furthermore, statistics on the shape of the syntax tree of the systems outputs (e.g., depth, breadth) could give some insight into the syntactic simplicity of outputs. This would improve the above evaluation protocol as its measure of simplicity is quite limited: ISiM, the main metric for this dimension, is oblivious to the adequacy of the simplification and its grammaticality.

It should be noted that automatic evaluation of ATS systems remains an open question [3, 1]. Most metrics fail to address some aspect of simplification altogether and may exhibit pitfalls in their evaluation of the aspects they do address. In particular high scores of the above metrics do not imply good simplification quality [3]. Furthermore, posing simplicity as solely a textual property without accounting for the competence of the intended audience is an inherently incomplete view of the problem.

6.6 ATS methods and results

6.6.1 Lexical simplification

Our modular system for lexical simplification (for English) requires no training data and allows us to fine-tune each module separately in order to improve the result. It consists of three modules operating consecutively: complex word identification (CWI), candidate generation, and candidate ranking.

Complex Word Identification. To identify the candidate for simplification, we first segment the input sentence into words before passing them through a vocabulary classifier (detailed in the paragraph about *Candidate Ranking*), along with the whole sentence as contextual information. The word with the highest complexity score becomes the candidate for simplification. In case of a tie, we prioritize

---

[7] https://github.com/feralvam/easse

verbs and select the word with the highest frequency in order to exclude inappropriate candidates for lexical simplification such as terminology.

Candidate Generation. To generate substitution candidates, we used the lexical substitution framework *LexSubGen* [4] and its best performing estimator *XLNet+emb* to generate 20 suggestions given a target complex word. We modified the post-processing of the original system to exclude the candidate lemmatization and get inflected suggestions, rather than lemmas. We kept the lowercase post-processor followed by target exclusion that uses lemmatization to detect and exclude all forms of the target word.

Candidate Ranking. We selected and ranked candidates based on a combination of their grammaticality, meaning preservation, and simplicity scores, through a simple heuristic giving twice as much weight to the simplicity score. The rank of each substitution w_n of the $N = 20$ generated candidates is determined as a function of its grammaticality $G \in \{0, 1\}$, simplicity $S \in [1, 6]$, and meaning preservation $M \in [1, N]$ scores, as presented in the following equation.

$$rank \quad w_{n \leq N} = G_{w_n} \times (S_{w_n} \times 2 + M_{w_n})$$

The top 10 candidates (or fewer) are those included in the submission.

To evaluate the **grammaticality** of a sentence given a substitute candidate, we compare the coarse-grained part-of-speech (POS) and morphological features of both complex word and candidate in context. We use spaCy[8] to tokenize and parse the sentence, making sure not to split hyphenated complex words, since LexSubGen does not support multi-word expressions. We assign a score of 1 to any candidate whose features (person, number, mood, tense, etc.) correspond to those of the target word and 0 otherwise.

To evaluate the effect of a substitution candidate on the **meaning** of the original sentence we compute the similarity of the two sentences as a sum of the cosine similarities between their token embeddings using BERTScore [28]. The higher the similarity between source and target sentences, the higher the chances that the substitution candidate meaning is close to the one of the complex word. Candidates are ranked by decreasing F1 score with the best candidate receiving a score of 1 and the last one a score equal to N.

We measure the **lexical complexity** of each candidate with a CEFR[9] vocabulary classifier trained on data from the English Vocabulary Profile[10] (EVP) [6]. EVP is a rich resource in British and American English that associates single words, phrasal verbs, phrases, and idioms not only with a CEFR level but with part of speech tags, definitions, dictionary examples, and examples from learner essays. The corpus also contains distinct entries for distinct meanings of polysemous words, each associated with its own difficulty level. For each substitution candidate we extract a semantic, contextual, dense vector representation from a pre-trained masked language model[11] [7] by first encoding the target word or multi-word expression (MWE) in context (using the dictionary and learner examples) and then aggregating all 12 hidden layers for all WordPieces. This representation of the dataset is then used to train a support vector classifier [18]. The resulting model is able to assign a difficulty level between 1 (A1) and 6 (C2) to the meaning of any word or MWE as determined by its context.

Results and discussion. We evaluated (partially)[12] our lexical simplification system (in English) on the TSAR-2022 dataset. The main results are reported in Table 6.2. The system outperforms the state-of-the-art *LSBert* baseline on 27 out of the total 51 metrics (including Precision and Recall).

---

[8] `https://spacy.io/` — v. 3.1.3 — en-core-web-lg

[9] The Common European Framework of Reference for Languages (CEFR) organizes language proficiency in six levels, A1 to C2.

[10] `https://www.englishprofile.org/american-english`

[11] https://huggingface.co/bert-base-uncased

[12] We did not evaluate the CWI part of the system and leveraged the fact that the TSAR-2022 dataset identifies complex words. The results thus presume a perfect complex word identification prior to candidate generation and ranking.

Table 6.2: Results of SOTA and our systems on selected metrics.

	ACC@1@top1	ACC@1	MAP@3	Potential@3	Precision@4	Recall@4
LSBert-baseline	0.303	0.597	0.407	0.823	0.412	0.193
Our system	0.236	0.544	0.382	0.831	0.418	0.196

The high Potential@3 suggests that the system would be useful to human editors whenever suggesting 3 (or more) candidates to choose from. The Accuracy@K@top1 of the system doubles when K = 3 as compared with K=1, which means that, 46% of the time, the most commonly suggested substitute is among our top 3 predictions. Table 6.3 contains a number of example simplifications for sentences from the TSAR-2022 corpus.

Table 6.3: Selected sentences with complex word (in bold) and the top candidate produced by our system.

Sentence with complex word	Top substitute
It **decomposes** to arsenic trioxide, elemental arsenic and iodine when heated in air at 200°C.	changes
Lebanon is sharply split along **sectarian** lines, with 18 religious sects.	religious
The stretch of DNA transcribed into an RNA molecule is called a transcription unit and **encodes** at least one gene.	codes
Obama earlier dropped from night skies into Kabul [...], **cementing** 10 years of U.S. aid for Afghanistan after NATO combat troops leave in 2014.	securing

6.6.2 Sentence simplification

Given the particulars of the problem and the needs of CBC/Radio-Canada, namely the lack of French-language resources, several approaches from the literature were considered. The selected approaches were UNTS [25], ACCESS [12], and MUSS [13]. Our preliminary results concern only the ALECTOR corpus and the BLEU and SARI metrics. UNTS was trained on the MLSUM and RFI datasets; ACCESS was trained on the translated Wikilarge corpus. MUSS was used as trained by its original authors [13]. The results are presented in Table 6.4.

Table 6.4: Results of selected models on the sentence-aligned French-language ALECTOR corpus.

Model	BLEU	SARI
UNTS	37.6	35.5
ACCESS	46.7	34.9
MUSS	38.9	38.1

The paraphrase-based approach, MUSS, achieves the highest SARI score, while ACCESS achieves the highest BLEU score but at the cost of a lower SARI. This is consistent with the reported state-of-the-art performance of MUSS on other datasets [13]. It should be noted, however, that these results are fragile in that ALECTOR only has a single reference output whereas both metrics are more reliable with several reference outputs [3, 27].

6.7 Conclusion

Evaluation for ATS tasks is not standard and the current metrics have known limitations [3]. In this work, a protocol to evaluate our ATS systems was identified. It comprises multiple metrics to capture the three dimensions of text simplification: grammaticality, meaning preservation, and simplicity. While imperfect, this metric ensemble allows for deeper analyses of our systems.

This work explored two avenues for text simplification: lexical simplification and end-to-end systems for sentence simplification. Our lexical simplification solution obtained performances similar to those of the shared task baseline [23] and showed promising results when considering the top 3 to 4 words. This suggests that it could be a good system to support expert annotators. This report also presents partial results for the ALECTOR dataset [10] with several end-to-end systems. While fragile due to the lack of multiple references, results of ACCESS [12] and MUSS [13], an ACCESS-based approach, suggest that further adjustments of this controllable model could further improve simplification quality.

To complete this work, a deeper evaluation of the presented approaches along the three ATS dimensions (meaning preservation, grammaticality, and simplicity) is needed to assess the quality of the discussed systems in an appropriate manner. Moreover this evaluation could be complemented by the use of translated multiple-reference corpora such as ASSET [1].

Bibliography

[1] Fernando Alva-Manchego, Louis Martin, Antoine Bordes, Carolina Scarton, Benoît Sagot, and Lucia Specia. ASSET: A dataset for tuning and evaluation of sentence simplification models with multiple rewriting transformations. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 4668–4679, Online, July 2020. Association for Computational Linguistics.

[2] Fernando Alva-Manchego, Louis Martin, Carolina Scarton, and Lucia Specia. EASSE: Easier automatic sentence simplification evaluation. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations, pages 49–54, Hong Kong, China, November 2019. Association for Computational Linguistics.

[3] Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. The (Un)Suitability of automatic evaluation metrics for text simplification. Comput. Linguist., 47(4):861–889, December 2021.

[4] Nikolay Arefyev, Boris Sheludko, Alexander Podolskiy, and Alexander Panchenko. Always keep your target in mind: Studying semantics and improving performance of neural lexical substitution. arXiv preprint arXiv:2206.11815, 2022.

[5] Mikel Artetxe and Holger Schwenk. Massively multilingual sentence embeddings for Zero-Shot Cross-Lingual transfer and beyond. Transactions of the Association for Computational Linguistics, 7:597–610, 2019.

[6] Annette Capel. Completing the English vocabulary profile: C1 and C2 vocabulary. English Profile Journal, 3, 2012.

[7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

[8] Yue Dong, Zichao Li, Mehdi Rezagholizadeh, and Jackie Chi Kit Cheung. EditNTS: An neural Programmer-Interpreter model for sentence simplification through explicit editing. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 3393–3402, Florence, Italy, July 2019. Association for Computational Linguistics.

[9] Cécile Fabre, Nabil Hathout, Lydia-Mai Ho-Dac, François Morlane-Hondère, Philippe Muller, Franck Sajous, Ludovic Tanguy, and Tim Van de Cruys. Présentation de l'atelier semdis 2014: sémantique distributionnelle pour la substitution lexicale et l'exploration de corpus spécialisés. In 21e Conférence sur le Traitement Automatique des Langues Naturelles (TALN 2014), pages 196–205, 2014.

[10] Núria Gala, Anaïs Tack, Ludivine Javourey-Drevet, Thomas François, and Johannes C Ziegler. Alector: A parallel corpus of simplified French texts with alignments of misreadings by poor and dyslexic readers. In Proceedings of the Twelfth Language Resources and Evaluation Conference, pages 1353–1361, Marseille, France, May 2020. European Language Resources Association.

[11] Reno Kriz, Marianna Apidianaki, and Chris Callison-Burch. Simple-QE: Better automatic quality estimation for text simplification. arXiv preprint arXiv:2012.12382, December 2020.

[12] Louis Martin, Éric de la Clergerie, Benoît Sagot, and Antoine Bordes. Controllable sentence simplification. In Proceedings of the Twelfth Language Resources and Evaluation Conference, pages 4689–4698, Marseille, France, May 2020. European Language Resources Association.

[13] Louis Martin, Angela Fan, Éric de la Clergerie, Antoine Bordes, and Benoît Sagot. MUSS: Multilingual unsupervised sentence simplification by mining paraphrases. arXiv preprint arXiv:2005.00352, May 2020.

[14] Lucas Mucida, Alcione Oliveira, and Maurilio Possi. Language-independent metric for measuring text simplification that does not require a parallel corpus. The International FLAIRS Conference Proceedings, 35, May 2022.

[15] Kostiantyn Omelianchuk, Vipul Raheja, and Oleksandr Skurzhanskyi. TExt SImplification by TAgging. arXiv preprint arXiv:2103.05070, 2021.

[16] Robert Östling, Yves Scherrer, Jörg Tiedemann, Gongbo Tang, and Tommi Nieminen. The Helsinki neural machine translation system. In Proceedings of the Second Conference on Machine Translation, pages 338–347. aclanthology.org, 2017.

[17] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic evaluation of machine translation. In Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02, pages 311–318, USA, July 2002. Association for Computational Linguistics.

[18] John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. Advances in large margin classifiers, 10(3):61–74, 1999.

[19] Sadaf Abdul Rauf, Anne-Laure Ligozat, François Yvon, Gabriel Illouz, and Thierry Hamon. Simplification automatique de texte dans un contexte de faibles ressources (Automatic text simplification: Approaching the problem in low-resource settings for French). In Actes de la 6e conférence conjointe Journées d'Études sur la Parole (JEP, 33e édition), Traitement Automatique des Langues Naturelles (TALN, 27e édition), Rencontre des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (RÉCITAL, 22e édition). Volume 2: Traitement Automatique des Langues Naturelles, pages 332–341. aclanthology.org, 2020.

[20] Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. MLSUM: The multilingual summarization corpus. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 8051–8067. aclanthology.org, 2020.

[21] Thomas Scialom, Louis Martin, Jacopo Staiano, Éric Villemonte de la Clergerie, and Benoît Sagot. Rethinking automatic evaluation in sentence simplification. arXiv preprint arXiv:2104.07560, April 2021.

[22] Sanja Stajner, Daniel Ferrés, Matthew Shardlow, Kai North, Marcos Zampieri, and Horacio Saggion. Lexical simplification benchmarks for English, Portuguese, and Spanish. arXiv preprint arXiv:2209.05301, 2022.

[23] Sanja Stajner, Maja Popovic, Horacio Saggion, Lucia Specia, and Mark Fishel. Shared task on quality assessment for text simplification. Training, 218(95):192, 2016.

[24] Elior Sulem, Omri Abend, and Ari Rappoport. Semantic structural evaluation for text simplification. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 685–696, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.

[25] Sai Surya, Abhijit Mishra, Anirban Laha, Parag Jain, and Karthik Sankaranarayanan. Unsupervised neural text simplification. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 2058–2068, Florence, Italy, July 2019. Association for Computational Linguistics.

[26] Sowmya Vajjala and Ivana Lucic. OneStopEnglish corpus: A new corpus for automatic readability assessment and text simplification. In Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications, pages 297–304, Stroudsburg, PA, USA, June 2018. Association for Computational Linguistics.

[27] Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. Optimizing statistical machine translation for text simplification. Transactions of the Association for Computational Linguistics, 4:401–415, 2016.

[28] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. BERTScore: Evaluating text generation with BERT. In Eighth International Conference on Learning Representations, April 2020.

[29] Xingxing Zhang and Mirella Lapata. Sentence simplification with deep reinforcement learning. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Stroudsburg, PA, USA, 2017. Association for Computational Linguistics.

7 Enumeration of maximal bicliques for helping detect at-risk companies

Karine Dufresne [a]
Mathieu Gervais-Dubé [b]
Nicolas Goulet [c]
Alain Hertz [b]
Odile Marcotte [d]
Hugues-Etienne Moisan-Plante [a]

[a] *Revenu Québec*

[b] *Département de mathématiques et de génie industriel, Polytechnique Montréal & GERAD*

[c] *Département d'informatique, UQAM*

[d] *UQAM & GERAD*

December 2022
Les Cahiers du GERAD

Copyright © 2022 GERAD, Dufresne, Gervais-Dubé, Goulet, Hertz, Marcotte, Moisan-Plante

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract: *Revenu Québec is implementing an algorithmic framework for detecting companies at risk of not complying with their fiscal duties. This framework requires (a) creating a bipartite graph in which the nodes are companies and resources and (b) computing all the maximal bicliques containing a given company, for each company in that graph. We present an algorithm for computing all the maximal bicliques in such a bipartite graph and discuss the results obtained by applying it to the graph supplied by Revenu Québec.*

7.1 Background

An important issue for governments across the world is the recovery of taxes owed by companies. Detecting companies at risk of fraud has until recently been a time-consuming process carried out by field auditors. Fraud detection algorithmic frameworks have been proposed to facilitate the work of these auditors and researchers at Revenu Québec have begun implementing such a framework, presented in articles by Van Vlasselaer et al. [1] and [2]. (Note that each article includes many references to articles on fraud detection, especially detection relying on the relationships between companies.) In what follows the phrase "at-risk company" means a "company at risk of fraud" (in the case of Belgium) or "at risk of tax non-compliance" (in the case of Revenu Québec).

To determine whether a company is at risk or not one can look at its intrinsic features or its network-based features. For instance, among the intrinsic features, an obvious one is the past involvement of a company in fraudulent activities. Network-based features are more subtle and involve the consideration of resources shared by two or more companies. For example, if two companies C and D have a common address (resp. supplier, employee, piece of equipment, etc.), this might point to a close link between these companies and therefore a high probability that C is at risk if D is (or vice versa). Of course more than two companies may share the same resource.

To formalize the concept of network-based feature we introduce the bipartite graph $G = (V_C, V_R, E)$: V_C is the set of companies, V_R the set of resources, and E the set of couples (x, y) where x is a company using resource y. In general a network-based feature will be computed using information extracted from the graph G, e.g., a parameter of G. In this article we focus on network-based features related to *bicliques* in G, where a biclique is defined as a complete bipartite subgraph of G induced by a set of nodes containing at least two companies and two resources. A maximal biclique is a biclique that is not contained in another biclique. Given a company x we can associate to x the number of maximal bicliques in G that contain x: this number is an example of network-based feature.

The researchers at Revenu Québec wish to compute the set of all maximal bicliques within a reasonable amount of time, which does not seem to be possible when one attempts to use the method outlined in Van Vlassselaer et al. [2]. Lu et al. [3] have proposed an algorithm for enumerating maximal bicliques in bipartite graphs and they cite several articles with applications of this problem to varied questions arising in the biological sciences. There are several definitions of the maximum biclique problem, of varied complexities (see for instance Peeters [4]), but the *enumeration* of all maximal bicliques is clearly a difficult problem (see Yang et al. [5]).

In the rest of this article we present the anonymized data set provided by Revenu Québec and the method used to solve the problem for this data set. We conclude by outlining avenues for future research.

7.2 The data set and our method

The data set provided by Revenu Québec included approximately 500,000 companies and 3.3 million resources, which it is not surprising since Québec has more than 8 million inhabitants. Enumerating all the maximal bicliques appears to be a daunting task but a quick look at the data set reveals that it contains a large number of nodes of degree 1, and more generally a large number of nodes of small degree. This observation led us to design an algorithm that looks at the nodes of the bipartite graph in increasing order of degrees.

First we note that a node of degree 1 cannot belong to a biclique, since we have assumed that a biclique contains at least two companies and two resources. Hence we can remove all nodes of degree 1 from the graph in an iterative manner, since the degrees of some nodes will decrease after the removal of the nodes of degree 1. When the graph G no longer contains any node of degree 1, we consider nodes of degree 2. If x has the two neighbors y and z, there is a maximal biclique containing x if and only if there is at least one node $w \neq x$ that is adjacent to y and z. If that condition is met, there is only one maximal biclique containing x, the one whose node set is $\{y, z\} \cup \{w | w \text{ is adjacent to } y \text{ and } z\}$. After considering node x we can remove it from G, whether it is contained in a maximal biclique or not. The removal of nodes of degree 2 may again "create" new nodes of degree 1 or 2 and these can be removed from G in a similar fashion.

We now turn to nodes of degree 3. If the three neighbours of x are denoted y, z, and w, there are four possibilities for a maximal biclique U containing x: (a) $y, z, w \in U$, (b) $y, z \in U, w \notin U$, (c) $y, w \in U, z \notin U$, and (d) $z, w \in U, y \notin U$. To each of these possibilities corresponds a potential maximal biclique: $B_1 = \{y, z, w\} \cup \{u | uy, uz, uw \in E\}$, $B_2 = \{y, z\} \cup \{u | uy, uz \in E\}$, $B_3 = \{y, w\} \cup \{u | uy, uw \in E\}$, and $B_4 = \{z, w\} \cup \{u | uz, uw \in E\}$. Note that we wrote "potential maximal biclique" because we have to make sure that B_i includes at least two companies and two resources and B_i is not included in another B_j for some j. We add all maximal bicliques found in this way to the list of maximal bicliques and remove x from G. In general, if N_x is the set of neighbours of x, where $|N_x|$ is at least 3, we need to consider all subsets of N_x of cardinality at least 2 and verify that the corresponding bicliques are indeed maximal.

We illustrate this process on the graph 7.1. First nodes 1 and 6 are removed from the graph; node a being then of degree 1, it is also removed from the graph and the new graph is depicted in Figure 7.2. Node f is considered next since it has degree 2: the only clique containing f is $\{d, e, f, 4, 5\}$ and node f is removed from the graph (yielding the graph 7.3). (We should also consider node 2 but will skip that step for the sake of illustrating the process on nodes of degree 3.) In Figure 7.3 we observe that node b is of degree 3 and we have to consider all subsets of $\{2, 3, 4\}$ with at least two nodes. The subset $\{2, 3, 4\}$ yields the maximal biclique $\{b, c, 2, 3, 4\}$, the subsets $\{2, 3\}$ and $\{2, 4\}$ yield respectively $\{b, c, 2, 3\}$ and $\{b, c, 2, 4\}$ (neither of which is maximal), and the subset $\{3, 4\}$ yields another maximal biclique, $\{b, c, d, e, 3, 4\}$.

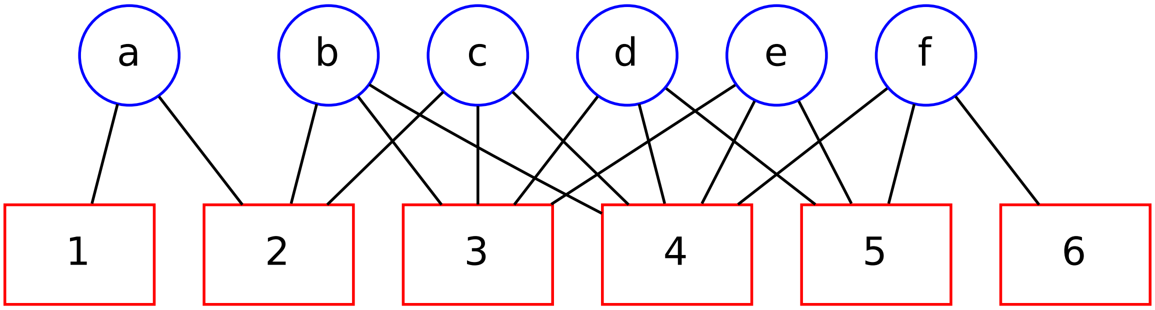

Figure 7.1: Initial graph.

As the degree increases this process consumes a lot of time: for instance, if u is a node of degree 4 (resp. 5) and $N(u)$ denotes the set of neighbours of u, there are 11 (resp. 26) subsets of $N(U)$ to consider. In the case of the anonymized data set provided by Revenu Québec, we remove iteratively all the nodes of degree at most 6, i.e., once a node (say, u) had been removed, we remove from the

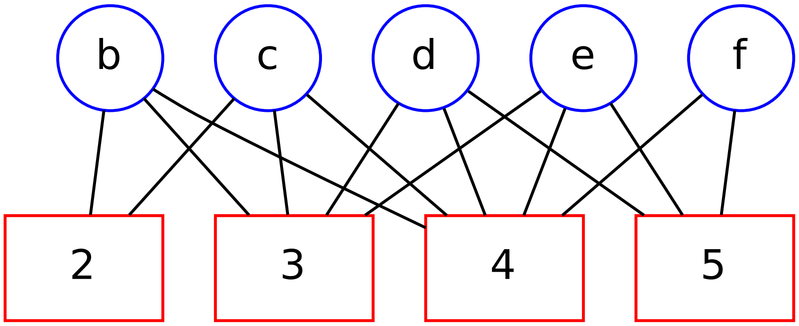

Figure 7.2: Graph after the removal of nodes of degree 1.

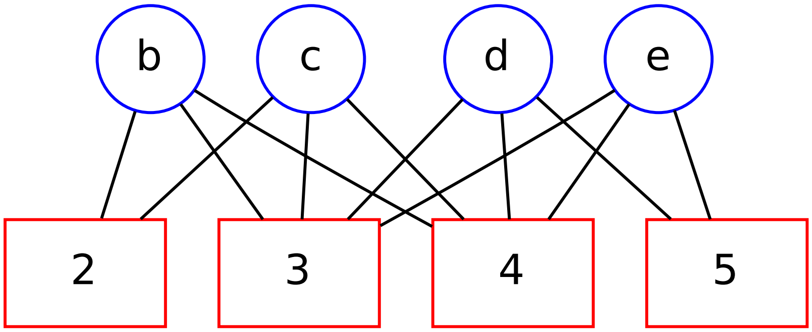

Figure 7.3: Graph after the removal of nodes of degre 1 and node f.

graph other nodes (for instance neighbours of u) whose current degree is at most 6. This yields a graph whose characteristics remind one of a *small-world network* (see Watts [6]): it has $10,075$ nodes and consists of one huge connected component (with 9922 nodes) and nine very small ones (all in the upper right-hand corner of Figure 7.4).

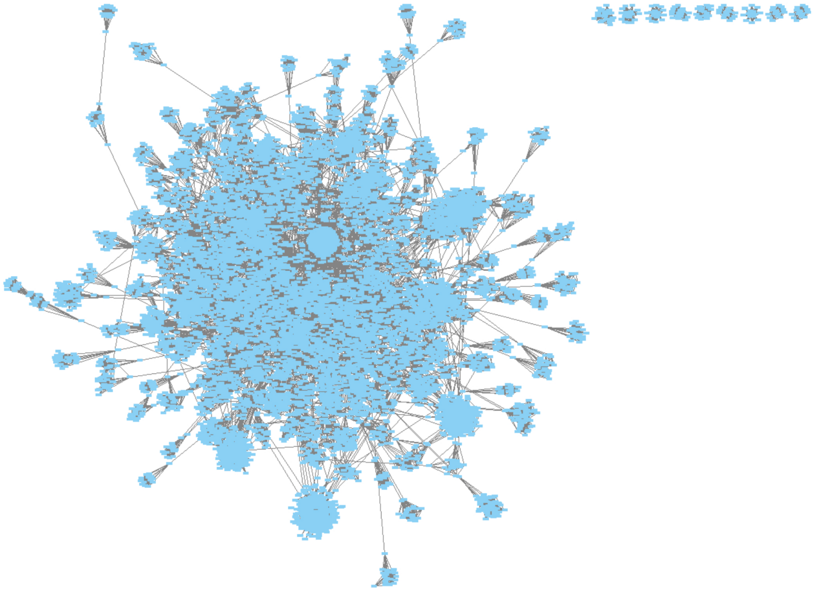

Figure 7.4: Graph of Revenu Québec after the removal of all nodes of degree 6 or less.

Table 7.1: Number of nodes as a function of the maximal degree of removed nodes.

Max degree	Nodes	Max degree	Nodes	Max degree	Nodes	Max degree	Nodes
1	1308308	8	4431	15	1125	46	926
2	614789	9	3028	16	1051	108	925
3	218306	10	2265	17	1000	125	924
4	64319	11	1984	18	973	132	403
5	21757	12	1685	19	929	136	366
6	10075	13	1580	20	928	138	0
7	6403	14	1408	21	927		

Maximal bicliques are computed throughout the scanning of nodes of degree at most 6, which consumes around 10 minutes of computer time. One can then apply the algorithm of Lu et al. [3] to the resulting graph of $10,075$ nodes and obtain (in around one minute and a half) a list of all candidate maximal bicliques in this graph. We use the phrase "candidate maximal bicliques" because a biclique B returned by the algorithm of Lu et al. may be contained in a biclique found while scanning the nodes of degree at most 6. To check that B is indeed maximal we consider every neighbour v of a node in B that was removed when scanning nodes of degree at most 6: if v is adjacent to every node in one of the parts of B, then B is not maximal. If no such v is found then B is maximal. We expect the checking to consume little time but did not have time to verify this during the workshop.

We also considered a "more agressive approach" in the first part of our method, i.e., we scanned iteratively all the nodes of degree at most 21. In that case the application of the algorithm of Lu et al. took only 9 milliseconds. Altogether reading the data, creating the graph in the software, scanning the nodes, and checking that the bicliques are indeed maximal took between 30 and 40 minutes. Therefore the scanning of nodes of degree at most 6 might be more efficient, provided the checking of the bicliques produced the algorithm of Lu et al. does not take too much time.

To conclude this section we present Table 7.1, describing the number of nodes as a function of the maximum degree of removed nodes. For instance, after removing iteratively all nodes of degree 1 from the original graph, the resulting graph has $1,308,308$ nodes. When the maximum degree of a removed node is 6 (resp. 21), the number of remaining nodes is $10,075$ (resp. 927). The last two columns records the data for degrees where a change occurs: for instance there is no change in the number of nodes between degree 21 and degree 45, while the number of nodes for degrees 136 and 137 is 366.

7.3 Future research

Finding maximal bicliques in bipartite graphs has many applications, in particular to problems arising in biology (see the list of references in the article by Lu et al., which includes a reference to a library of graphs). During the workshop our investigations have shown that graph sparsity is very helpful for enumerating maximal bicliques, as indeed for solving many graph optimization problems. We intend to pursue this research avenue by examining several graph classes and proposing algorithms for them.

Bibliography

[1] Véronique Van Vlasselaer, Tina Eliassi-Rad, Leman Akoglu, Monique Snoeck, Bart Baesens (2016) GOTCHA! Network-Based Fraud Detection for Social Security Fraud. Management Science 63(9):3090–3110

[2] V. Van Vlasselaer, L. Akoglu, T. Eliassi-Rad, M. Snoeck and B. Baesens, Guilt-by-Constellation: Fraud Detection by Suspicious Clique Memberships, 2015 48th Hawaii International Conference on System Sciences, 2015, pp. 918–927, doi: 10.1109/HICSS.2015.114

[3] Lu, Yuping, Phillips, Charles A., Langston, Michael A., 2020 Biclique: an R package for maximal biclique enumeration in bipartite graphs BMC Research Notes 13 https://doi.org/10.1186/s13104-020-04955-0

[4] Peeters, MJP 2000, The Maximum Edge Biclique Problem is NP-Complete. FEW Research Memorandum, vol. 789, Operations research, Tilburg.

[5] Jianye Yang, Yun Peng, Wenjie Zhang (p,q)-biclique Counting and Enumeration for Large Sparse Bipartite Graphs Proceedings of the VLDB Endowment, 2021–2022, 15(2):141–153

[6] Watts, D. J.; Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks (PDF). Nature. 393 (6684):440–442.