Trivella, Alessio; Wang, Pengling; Corman, Francesco

Article

The impact of wind on energy-efficient train control

EURO Journal on Transportation and Logistics (EJTL)

Provided in Cooperation with:
Association of European Operational Research Societies (EURO), Fribourg

Suggested Citation: Trivella, Alessio; Wang, Pengling; Corman, Francesco (2021) : The impact of wind on energy-efficient train control, EURO Journal on Transportation and Logistics (EJTL), ISSN 2192-4384, Elsevier, Amsterdam, Vol. 10, Iss. 1, pp. 1-13, https://doi.org/10.1016/j.ejtl.2020.100013

EURO Journal on Transportation and Logistics 10 (2021) 100013

Contents lists available at ScienceDirect

EURO Journal on Transportation and Logistics

journal homepage: www.journals.elsevier.com/euro-journal-on-transportation-and-logistics

The impact of wind on energy-efficient train control

Alessio Trivella *, Pengling Wang, Francesco Corman

Institute for Transport Planning and Systems, ETH Zürich, 8093, Zurich, Switzerland

ARTICLE INFO

Keywords:
Train trajectory optimization
Energy-efficient train operations
Dynamic programming
Train resistances
Wind variability

ABSTRACT

An energy-efficient train trajectory corresponds to the speed profile of a train between two stations that minimizes energy consumption while respecting the scheduled arrival time and operational constraints such as speed limits. Determining this trajectory is a well-known problem in the operations research and transport literature, but has so far been studied without accounting for stochastic variables like weather conditions or train load that in reality vary in each journey. These variables have an impact on the train resistance, which in turn affects the energy consumption. In this paper, we focus on wind variability and propose a train resistance equation that accounts for the impact of wind speed and direction explicitly on the train motion. Based on this equation, we compute the energy-efficient speed profile that exploits the knowledge of wind available before train departure, i.e., wind measurements and forecasts. Specifically, we: (i) construct a distance-speed network that relies on a new non-linear discretization of speed values and embeds the physical train motion relations updated with the wind data, and (ii) compute the energy-efficient trajectory by combining a line-search framework with a dynamic programming shortest path algorithm. Extensive numerical experiments reveal that our "wind-aware" train trajectories present different shape and reduce energy consumption compared to traditional speed profiles computed regardless of any wind information.

1. Introduction

The global transport emissions have kept increasing over the last 20 years, with the transport sector accounting for 24% of direct CO_2 emissions from fuel combustion (IEA, 2019). Despite railways being one of the most energy-efficient transport modes, reducing energy consumption is considered as a major challenge and priority in modern railway transportation. In fact, the energy consumed by trains represents one of the largest costs incurred by railway operators. Moreover, there is large pressure on operators due to societal and environmental targets to increase the sustainability of railway transportation by improving efficiency and cutting emissions (Wang et al., 2011; Railenergy, 2016; UIC, 2019). Consequently, energy efficiency has emerged as a prominent subject within railway operations in both industry and academia. For example, 28 European railway operators have committed to reduce CO_2 emissions per passenger kilometer and per tonne kilometer by 50% by 2030 (UIC, 2019). These efforts are also signaled by public web pages of many railway operators (see, e.g., FS, 2019, NS, 2019; and SBB, 2019) and by the increasing amount of scientific literature on the topic (Liu and Golovitcher, 2003; Gunselmann, 2005; Howlett and Pudney, 2012; González-Gil et al., 2014; Lu et al., 2014; Yang et al., 2016; Scheepmaker et al., 2017; De Martinis and Corman, 2018). In addition to improvements in vehicle technology (Gunselmann, 2005), infrastructure and building design (González-Gil et al., 2015), and loading of freight trains (Lai et al., 2008), one important option to improve energy-efficiency in railways is adopting driving strategies that minimize the energy consumption of individual trains while satisfying the timetable and additional operational constraints such as speed limits. A driving strategy is typically represented in the form of a speed/distance profile and will be called a *train trajectory* or simply a speed profile hereafter. Determining an energy-efficient train trajectory can lead to a potential energy saving in the range 5–20% (Hansen and Pachl, 2014), hence, it is an attractive option for railway companies to reduce energy consumption as it does not require infrastructural changes or major investments. In practical train control, enforcing these speed profiles is possible with the support of automated driver advice systems (DAS), which are being implemented by a growing number of railways around the world. A DAS can indeed monitor the train movement and provide updated information to the driver about target arrival time as well as control advices to achieve this target using energy-efficient driving (Panou et al., 2013).

In the scientific literature, computing energy-efficient speed profiles is usually referred to as the *train trajectory optimization problem*. In its

* Corresponding author.
E-mail addresses: alessio.trivella@ivt.baug.ethz.ch, alessio.trivella@alice.it (A. Trivella).

https://doi.org/10.1016/j.ejtl.2020.100013
Received 9 January 2020; Received in revised form 12 March 2020; Accepted 2 June 2020

simplest setting in which a train runs on a flat track under a single speed limit, the optimal trajectory can be characterized analytically as a sequence of four driving regimes: (1) maximum acceleration, (2) speed holding, also called cruising, (3) coasting, and (4) maximum braking, in this order. However, in more realistic settings involving varying speed limits, slopes, and possibly other constraints, analytical derivations are generally not available and finding the optimal train trajectory is a challenging problem (Albrecht et al., 2016a, b). Therefore, a multitude of solution approaches have been developed in the literature, which, at high level, can be grouped into three categories:

1. The “indirect” methods formulate the problem as a continuous boundary value problem using differential equations (Cheng and Howlett, 1992; Howlett, 2000; Liu and Golovitcher, 2003; Howlett and Pudney, 2012; Albrecht et al., 2016a, b).
2. The “direct” methods formulate and solve the problem by means of mathematical programming techniques (usually mixed-integer or non-linear; Wang et al., 2013; Wang and Goverde, 2016a, b; Ye and Liu, 2017; Luan et al., 2018a, b).
3. The dynamic programming (DP) approaches rely on constructing a network using discrete distance, time, or speed quantities, and executing indeed a dynamic programming algorithm (Franke et al., 2000; Ko et al., 2004; Lu et al., 2013; Ghaviha et al., 2017; Haahr et al., 2017; Zhou et al., 2017). These papers mainly differ regarding the specific constraints considered and the way the underlying network is constructed. For instance Haahr et al. (2017) rely on a distance-speed network while Zhou et al. (2017) use a distance-speed-time network. The solution methodology that we develop in this paper adds to this line of work.

The extant literature we are aware of has studied the train trajectory optimization problem in an environment which is independent from the specific journey-dependent conditions. In other words, speed profiles are computed based on the timetable and available running time, i.e., static information published in advance. In reality, there are factors such as wind, rain, line voltage, train load, and resistances, that vary across journeys on the same track and may affect energy consumption and the optimal driving strategy (Powell and Palacín, 2015; De Martinis and Corman, 2018; Wang et al., 2020). Our intuition is that if we were able to measure or forecast these factors and somehow account for them *proactively* before computing a train trajectory, then the optimal trajectory would potentially be more energy-efficient because it would adapt to more information compared to the timetable alone. For example, one would expect the shape of an optimal train trajectory to be altered in the presence of strong winds, e.g., by holding a lower speed or by starting coasting earlier or later. Thus, our research questions is the following: Can we exploit the short term knowledge of journey-dependent variables, such as weather data available before train departure, to proactively compute more energy-efficient speed profiles?

To address this question, in this paper we focus on the example of wind and propose a train resistance model that explicitly accounts for the effect of wind speed and direction on the moving train. In general, the train motion is governed by well-established physical equations (see e.g., Hansen and Pachl, 2014, or Wang and Rakha, 2018 for a recent study on train dynamics) and there exist simple models from the academic and practitioner literature to describe the effect of aerodynamic drag on railway vehicles (Peters, 1990; Schetz, 2001; Tian, 2009). This paper builds on this literature by proposing a more detailed and flexible wind resistance model that decomposes the impact of wind speed and direction with respect to the train into a traversal and a longitudinal component. We show that our model is consistent with the simpler models from the literature in case of pure headwind or tailwind, i.e., in case there is no cross-wind. Although additional empirical research may be needed to validate this model, our focus is on studying the impact of wind on train resistance in the context of train trajectory optimization, which has not been done before to our knowledge.

An optimal (energy-efficient) train trajectory that exploits wind information is henceforth referred to as a *wind-aware* trajectory. While determining this trajectory we assume that wind speed and direction are function of the location and known in advance (before departure), and that the train does not recover energy while braking. Finding this trajectory requires solving a bi-objective problem with objectives being the energy consumption and the travel time. We tackle this problem by proposing a method that combines a simple line search approach with a DP shortest path algorithm in a distance-speed network with arc costs adjusted for the wind state. Our distance-speed network relies on a novel non-linear discretization of speed values that results in a more balanced network and might therefore improve the quality of the trajectories. Moreover, despite the fact that dynamic-programming based solutions exist for train trajectory optimization, our algorithm that combines solving shortest paths with line search is new and computationally very efficient. From a multi-objective optimization theory perspective, our method is an application of the weighted sum scalarization technique, which is not guaranteed in general to reach all Pareto-optimal solutions (Boland et al., 2015; Ehrgott, 2006). However, we show by experimentation that this method is suitable for solving the train trajectory optimization problem by approximating the energy-time Pareto frontier and showing that it is very dense of points, i.e., it shows no big holes that would hint the possibility of unreachable points due to non-convexities.

We perform a numerical study with realistic train data and different wind scenarios, and show that our wind-aware trajectories deviate from classical trajectories computed independent of wind and can be more energy-efficient. Specifically, our wind-aware trajectories lead to an energy saving that: (i) is related to the wind speed and impact angle, and (ii) can be up to 10% under strong headwinds and about 0–2% under low winds. By analyzing the shape of wind-aware trajectories, we derive insights on how the speed profile changes depending on the wind. In particular, we notice that the coasting phases can vary significantly in both shape and starting location. With our results, we show that there may be a practical benefit in including information (measurements, forecasts) about external factors such as wind speed and direction in the computation of train trajectories. Our next step could be to extend this model to a pervasive description of uncontrollable externals factors in railway networks (like people onboard, friction coefficients, energy conversion efficiency, etc.), allowing to classify their relative effect and opportunities.

To summarize, the contributions of this paper to the train trajectory optimization and railway operations literature are the following: (i) we extend the simple train resistance equations from the literature to incorporate more flexibly the effect of wind speed and direction on train motion, (ii) we introduce a new non-linear discretization of speed values resulting in more balanced distance-speed networks, (iii) we develop a simple yet powerful optimization scheme that allows to obtain energy-efficient trajectories in seconds by combining line search with dynamic programming, and (iv) we show that accounting for wind data (or in general, journey-dependent conditions, uncontrollable conditions that can be known with limited knowledge at some specific time only) when optimizing speed profiles might have a significant impact on the train energy consumption.

The rest of the paper is structured as follows. In Section 2, we review the classic train dynamics and present our train resistance equation that embeds the effect of wind. In Section 3, we describe our approach to build the distance-speed network and our method to optimize the train speed profile. In Section 4, we introduce the setup of our numerical experiments and present our results. We conclude in Section 5 and provide directions for future research.

2. Wind effect on train motion

In Section 2.1, we present the standard physical model of train dynamics. In Section 2.2, we introduce a new train resistance equation that is grounded on the literature on aerodynamic drag and incorporates the effect of wind speed and direction. Units are reported in square brackets.

2.1. Classical train dynamics

The motion of a railway vehicle is governed by the following well-known differential equations (Hansen and Pachl, 2014):

$$\frac{dv(s)}{ds} = \frac{f(s) - R_{line}(s) - R_{train}(v)}{\rho \cdot m \cdot v(s)}; \tag{1a}$$

$$\frac{dt(s)}{ds} = \frac{1}{v(s)}, \tag{1b}$$

where s is the distance traversed by the vehicle [m], $v(s)$ is the vehicle speed [m/s], $t(s)$ is the time [s], $f(s)$ represents the traction force when positive and braking force when negative [N], and $R_{line}(s)$ and $R_{train}(v)$ are the line and train resistances [N], respectively. The parameter m is the train mass [kg] and ρ the rotating mass factor [−].

The line resistance term $R_{line}(s)$ in (1a) is caused by curve resistance, tunnel resistance, and gravitational pull due to track slopes (Wang et al., 2011; Vardy and Reinke, 1999). The gravitational pull is usually considered as the predominant component in $R_{line}(s)$ (Pachl, 2002) and is computed as $m\,g\,\sin(q(s))$, where $g \simeq 9.8\ \text{m/s}^2$ is the gravitational acceleration constant and $q(s)$ [radians] is the slope angle as a function of s. Many papers accounting for line resistance in train trajectory optimization consider indeed the gravitational pull component alone (Haahr et al., 2017; Wang and Goverde, 2016a). We refer to Pachl (2002) and Wang et al. (2011) for more detailed discussions on the other line resistance components.

The train resistance $R_{train}(v)$, is composed by frictional, mechanical, and air resistances and is commonly expressed as a quadratic function of train speed (Rochard and Schmid, 2000; Pachl, 2002; Tian, 2009; Wang et al., 2011; Hansen and Pachl, 2014):

$$R_{train}(v) = c_1 + c_2\ v + c_3\ v^2, \tag{2}$$

where $c_1 \geq 0$, $c_2 \geq 0$, and $c_3 \geq 0$ are vehicle-dependent constant scalars with unit [N], [N (s/m)], and [N (s/m)^2], respectively. Based on this general formula, several railway companies have established different approximations to model resistances in their own types of trains (Pachl, 2002). As discussed in Section 1, the train resistance in practice is not only a function of speed but is also affected by external variables like wind and rain that change from journey to journey, even when considering the same train running between the same pair of stations. Therefore, in the next section we will discuss how to incorporate some of these variables into Equation (2) to obtain a more precise and journey-dependent description of the actual resistance.

In addition to Equation (1), the train movement is constrained by upper bounds on the power generated by the engine $f(s)\ v(s) \leq p^{MAX}$ [W], on the traction force $f(s) \leq f^{MAX}$, and on the speed $v(s) \leq v^{MAX}(s)$, where $v^{MAX}(s)$ is the speed limit enforced at location s. We use distance instead of time as independent variable in Equation (1) because the speed limit is indeed a function of the location s, as is the line resistance. Moreover, the acceleration is bounded both from below and from above by $a^{MIN} \leq dv(s)/dt(s) \leq a^{MAX}$ to ensure sufficient riding comfort to passengers. Notice that a^{MIN} is negative and corresponds to the maximum deceleration. Time windows at passage points could be enforced as well (Wang and Goverde, 2016a).

Finally, the energy consumed between any two locations s_1 and s_2 can be computed as the integral of the force exerted by the train in this interval:

$$E_{s_1}^{s_2} = \int_{s_1}^{s_2} \max\{f(s),\ 0\}\ ds. \tag{3}$$

Notice that, given a constant and uniform train acceleration $a = a(s) = dv(s)/ds$ between two locations s_1 and s_2, the energy consumption varies as a function of the train resistance. In fact, as can be seen in Equation (1a), when the resistance increases, the force $f(s)$ needed to hold a constant acceleration has to increase too. Thus, the resulting energy consumption increases as a result of Equation (3). As a special case, this consideration is also true when the acceleration is equal to zero, i.e., the train is holding a constant speed.

2.2. A new train resistance equation

In this section, we focus on the impact of wind on the resistance of a train and propose a formal description of this relation. We denote by $\vec{v}$ the vector representing the train speed relative to the track/ground and by $\vec{w}$ the real wind vector, i.e., the wind speed relative to the ground. These vectors have length v and w, respectively. The angle between them $\theta \in [0, 2\pi]$ [radians] is called the *impact angle* of wind. Fig. 1 provides a visual representation of $\vec{v}$, $\vec{w}$, and θ.

Next, we denote by $\vec{v}_A$ the vector representing the speed of the train relative to the air. This vector can be defined as the vector difference of $\vec{v}$ and $\vec{w}$, that is, $\vec{v}_A := \vec{v} - \vec{w}$. Its opposite $\vec{w}_A = -\vec{v}_A$ represents the relative wind speed as experienced by the train in motion and is usually referred to as the *apparent wind*, while the angle γ between $\vec{v}$ and $\vec{v}_A$ is known as the *yaw angle* (Schetz, 2001; Bomhauer-Beins, 2019). We illustrate vectors $\vec{v}_A$ and $\vec{w}_A$ and the resulting yaw angle γ in Fig. 2.

Following Bomhauer-Beins (2019), we can express γ and the magnitude of $\vec{v}_A$, i.e. $v_A = w_A$, as a function of v, w, and θ based on simple trigonometric relations, as follows.

$$v_A = v_A(v, w, \theta) = \sqrt{(v - w\cos\theta)^2 + (w\sin\theta)^2}, \tag{4a}$$

$$\gamma = \gamma(v, w, \theta) = \arcsin\left(\frac{w\sin\theta}{v_A}\right). \tag{4b}$$

We relegate the details of this simple derivation to the appendix.

It is commonly assumed in the literature that wind affects the quadratic term (i.e., the air resistance, or aerodynamic drag) in Equation (2) and that the impact of wind on a moving train depends on the two parameters v_A and γ (Peters, 1990; Schetz, 2001; Tian, 2009), which ultimately are functions of v, w, and θ as shown by Equation (4). A simple model in the literature (Peters, 1990; Schetz, 2001) that does not account for cross-winds, i.e., it is only applicable to tailwinds or headwinds parallel to the train direction ($\theta = 0$ or $\theta = \pi$), states that

$$R_{train}(v, v_A) = R_{train}(v, w, \theta) = c_1 + c_2\ v + c_3\ v_A^2. \tag{5}$$

The only difference between (2) and (5) is that v in the quadratic component of the former equation has been replaced in the latter equation by v_A, which is equal to $v_A = v - w$ if $\theta = 0$ (tailwind) or $v_A = v + w$ if $\theta = \pi$ (headwind). Although Equation (5) is generally acknowledged by academics and practitioners, literature that considers the effect of cross-winds on train resistance is quite scarce. Peters (1990) and Schetz (2001) model cross-winds by making the drag coefficient c_3 dependent on the yaw angle γ as $\tilde{c}_3(\gamma) = c_3 \cdot (1 + 0.02|\overline{\gamma}|)$, where $\overline{\gamma}$ is the measure in degrees of γ and c_3 is the same drag coefficient that appears in (2) and (5), i.e., the

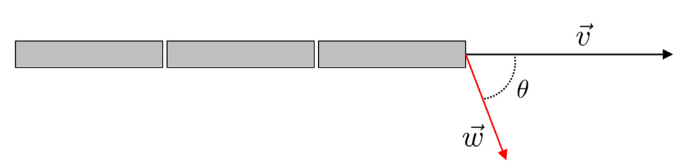

Fig. 1. Stylized train viewed from above with an illustration of train speed vector (black arrow), wind speed vector (red arrow), and impact angle θ. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

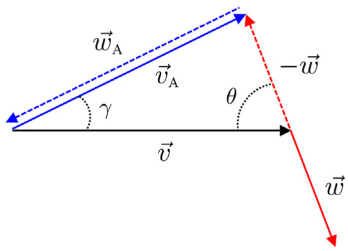

Fig. 2. Illustration of $\vec{v}_A$, $\vec{w}_A$, and γ.

baseline coefficient without cross-wind.

In this paper, we follow in spirit the approach of the aforementioned papers by making the drag coefficient c_3 wind-dependent. However, we propose a richer model for this coefficient that factors the impact of the apparent wind on the longitudinal and traversal surfaces of the train. To do this, we start by approximating the train by a parallelepiped with length L, width l, and height h. This means that the lateral sides of the train are described by rectangles with dimensions $L\times h$ whereas the front and back surfaces by $l \times h$ rectangles. Assume $\gamma \neq 0$, that is, the train is exposed to cross-wind, and consider the lateral side of the train which is hit by wind through the apparent wind vector $\vec{w}_A$. The surface that this vector actually sees is determined by projecting the lateral rectangle onto a plane orthogonal to $\vec{w}_A$, resulting in a smaller (or equal) rectangle with area $L\,h\,|\sin\gamma|$. In particular, when the angle γ is close to zero, then $L\,h\,|\sin\gamma|$ is also close to zero. In contrast, when γ is close to $\pi/2$, then $L\,h\,|\sin\gamma|$ is close to the original area $L\,h$. To better illustrate this projection, we use the examples in Fig. 3 to show the lateral area of the train and its projections under yaw angles of 20 and 60°. For instance, if $L = 100$ m and $h = 4$ m, i.e. the lateral train area is 400 m^2, then the area of the resulting projected rectangle is approximately 137 m^2 and 346 m^2, respectively, for the examples in Fig. 3(a) and (b).

Similarly, to determine the wind effect on the longitudinal train component, we consider the surface among front and back that is hit by $\vec{w}_A$, and project it onto a plane orthogonal to the direction of this vector. The resulting rectangle has an area of $l\,h\,|\cos\gamma|$. We use $\cos(\gamma)$ instead of $\sin(\gamma)$ here simply because the orientation of the front/back and lateral surface differs by $\pi/2$.

We now use the projected train surfaces to define a wind-dependent drag coefficient

$$c'_3(\gamma) := \frac{c_3}{h\,l}(\xi_1\,h\,l\,|\cos\gamma| + \xi_2\,h\,L\,|\sin\gamma|), \tag{6}$$

where ξ_1 and ξ_2 are positive parameters [$-$]. To build some intuition, we can think of (6) as the sum of a traversal and a longitudinal wind effect. The traversal component is related to the force that wind exerts on the lateral side of the train, hence it is proportional to the projected area $L\,h\,|\sin\gamma|$ orthogonal to the apparent wind. Similarly, the longitudinal component is related to the force that wind exerts on the front/back of the train and is proportional to $l\,h\,|\cos\gamma|$. The parameters ξ_1 and ξ_2 are introduced in (6) to provide a more flexible balance between the traversal and longitudinal component of the effect and can also be used to model the specific train shape. For instance, the improved aerodynamic design of high speed trains can be associated with lower values for ξ_1 and ξ_2. Our final expression of train resistance is defined by

$$R'_{train}(v, v_A, \gamma) = R'_{train}(v, w, \theta) = c_1 + c_2\,v + c'_3(\gamma)\,v_A^2. \tag{7}$$

In the special case in which $\xi_1 = 1$ and $\gamma = 0$, i.e. there is no cross-wind, in (6) we have $|\sin\gamma| = 0$, $|\cos\gamma| = 1$ and hence $c'_3(\gamma) = c_3$ thanks to the scaling coefficient $c_3/(h\,l)$. Consequently, (7) reduces to (5), that is, our model is consistent with and extends the existing model in the literature.

To illustrate, in Fig. 4(a) and (b) we display, respectively, the value of v_A and γ depending on the wind impact angle $\theta \in [0, 2\pi]$, given $v = 20$ m/s and $w = 10$ m/s. Moreover, in Fig. 4(c) we show the corresponding resistance values $R'_{train}(v, w, \theta)$ when further assuming the following parameters: $L = 138$ m, $l = 3.2$ m, $h = 4.4$ m, $c_1 = 5.8$ kN, $c_2 = 0.072$ kN s/m, $c_3 = 0.013$ kN $(\text{s/m})^2$, $\xi_1 = 1$, and three different values for $\xi_2 \in \{0, 0.025, 0.05\}$. Clearly, the relative effect of traversal and longitudinal impact depends on the value of ξ_1 and ξ_2. As shown in Fig. 4(c), increasing ξ_2 does not affect the resistance when $\theta = 0$ or $\theta = \pi$ but it does affect the resistance for other wind directions that involve cross-wind. Finding the actual values for ξ_1 and ξ_2 would require calibrating (6) and (7) from empirical measurements, which are unavailable in this study. Intuitively, one would expect a value for ξ_1 that is closer to 1 and a much smaller value for ξ_2. In our numerical study in Section 4, we perform a sensitivity analysis where we test different values for these parameters.

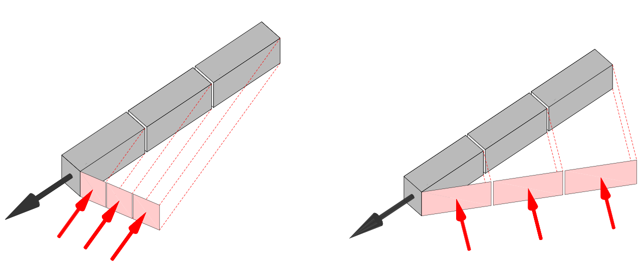

(a) $\gamma = 0.3491$ (20 degrees). (b) $\gamma = 1.047$ (60 degrees).

Fig. 3. Stylized train and projected lateral area under different yaw angles. The black and red arrows represent respectively the true train speed $\vec{v}$ and the apparent wind $\vec{w}_A$. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

3. Trajectory optimization including wind

In this section, we discuss our methodology to optimize train trajectories in the presence of wind. As mentioned in Section 1, computing the energy-efficient optimal trajectory in continuous time, distance, and speed is hard (regardless of wind) and can be done analytically only in very simple settings. To overcome this challenge, in this paper we develop an algorithm that relies on discrete distance and speed quantities and falls in the class of dynamic programming. In particular, in Section 3.1, we construct a distance-speed network that incorporates wind data and is based on a new non-linear discretization of speed values. In Section 3.2, we present a speed optimization method that combines a dynamic programming algorithm with line search. Although the network we define in Section 3.1 uses wind information, the solution method that we develop in Section 3.2 is not specifically wind-related and can be used to solve the standard train trajectory optimization problem.

3.1. Distance-speed network

We define our distance-speed network as a directed acyclic graph $\mathscr{G} = (\mathscr{N}, \mathscr{A})$, where $\mathscr{N}$ is the set of nodes and $\mathscr{A}$ the set of arcs (i.e., directed edges). Nodes $n \in \mathscr{N}$ and arcs $b \in \mathscr{A}$ in this graph represent, respectively, distance-speed pairs and feasible train movements between consecutive locations. We describe these sets in details in the following.

3.1.1. Node set

To construct $\mathscr{N}$, we start by discretizing distance and defining a finite set of $I + 1$ locations named $\mathscr{S} = \{s_0 \leq \ldots \leq s_i \leq \ldots \leq s_I\}$, where $s_0 = 0$ corresponds to the departure station and $s_I = S$ to the arrival station (S denotes the total distance between the two stations). In our numerical experiments we use a uniform discretization of distance, meaning that $s_{i+1} - s_i = \Delta s = S/I$ for all $i = 1, \ldots, I$. However, the method that we develop in Section 3.2 to optimize trajectories works independently of the discretization used.

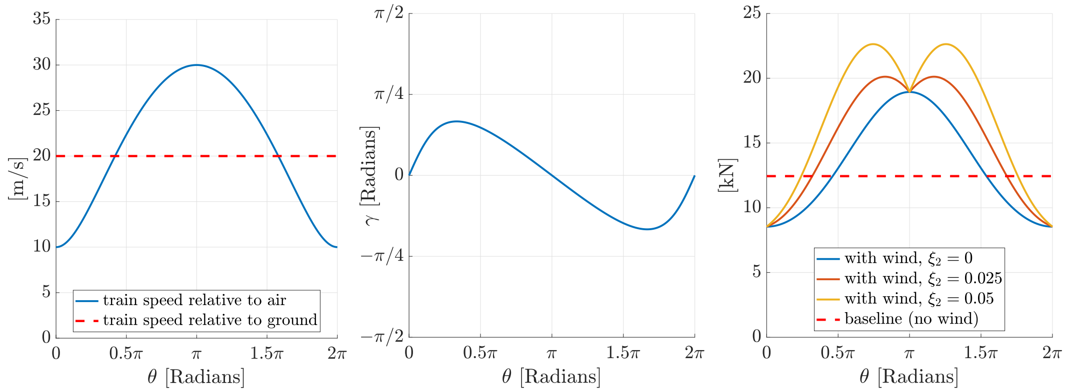

(a) Relative train speed. (b) Yaw angle. (c) Train resistance.

Fig. 4. Illustration of relative train speed v_A, yaw angle γ, and train resistance R'_{train} as functions of wind impact angle θ for fixed values of v and w.

To discretize speed, we consider the maximum speed limit over the entire track, which is $V^{MAX} = \max\{v^{MAX}(s), s \in [0, S]\}$, and define a set of $J+1$ speed points $\mathscr{V} = \{v_0 \leq \ldots \leq v_j \leq \ldots \leq v_J\}$, where $v_0 = 0$ and $v_J = V^{MAX}$. Again, we could specify these points using a uniform discretization (the methodology is not affected by the speed discretization choice). We notice however that, given two consecutive locations s_i and s_{i+1} and an initial speed v_j at location s_i, the range of speed values v_h that can be reached at s_{i+1} varies substantially depending on the initial speed v_j.

To elaborate, consider the example in Fig. 5 where the interval Δs is equal to 100 m, minimum and maximum accelerations are $a^{MIN} = -0.8$ m/s^2 and $a^{MAX} = 0.6$ m/s^2, and the track slope is zero.

If the initial speed is $v_j = 40$ km/h, then the train is able to speed up to 56 km/h or to decrease speed to 0 km/h, i.e. to stop, before arriving at s_{i+1}. Instead, if the initial speed is $v_j = 100$ km/h, then in the same distance the train can only increase or decrease speed, respectively, by 7 and 11 km/h. Therefore, the range of reachable speed values in km/h is $[0, 56]$ in the first case, while in the second case this range is just $[89, 107]$, i.e., it is more than three times smaller than the former interval. This observation is somehow intuitive because the higher the speed, the shorter the time needed for the train to run a given distance, the smaller the flexibility to change speed. As a consequence, choosing a uniform speed discretization entails a number of reachable speed points that quickly decreases as the speed increases, resulting in limited flexibility to adjust the train trajectory in the portions of the distance-speed network where the speed is high. To resolve this issue, we make use of a non-linear discretization based on a simple power function

$$v_j := (j/J)^{\varphi} \cdot V^{MAX}, \ \forall j = 0, \ldots, J, \tag{8}$$

with $\varphi \in (0, 1]$. When $\varphi = 1$, Equation (8) corresponds to the uniform discretization. As we reduce φ, the density of our $J+1$ points increases for high values of speed and decreases for low values of speed. In other words, decreasing φ results in speed points that concentrate more and more towards V^{MAX}. For instance, the parameter φ could be set such that a similar number of reachable points is obtained from any initial speed v_j, resulting in a network that is more balanced in the sense that the number of feasible controls is similar at any node. To the best of our knowledge, no previous research has employed non-linear discretization techniques in train trajectory optimization.

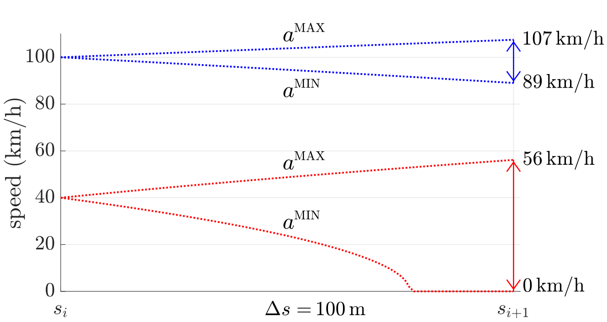

Fig. 5. Range of reachable speed values between two consecutive discretized locations.

Recall that speed limits are location-dependent. Once the set $\mathscr{V}$ is constructed according to Equation (8), we define location-dependent speed-point sets $\mathscr{V}_i$, for $i = 0, \ldots, I$, each being the subset of $\mathscr{V}$ that satisfies the respective speed restriction at location s_i. Formally, $\mathscr{V}_i = \{v_j \in \mathscr{V} : v_j \leq v^{MAX}(s_i)\} \cup v^{MAX}(s_i)$. We added $v^{MAX}(s_i)$ to this set to allow the train to reach (and hold) precisely the maximum permitted speed. Thus, we can define the nodes $n \in \mathscr{N}$ in our network as distance-speed pairs $n = (s_i, v_j) \in \mathscr{S} \times \mathscr{V}_i$.

Finally, notice that some of the nodes close to the points where the speed limits change will never be reached by feasible trajectories. To illustrate, consider the portion of a distance-speed graph in Fig. 6, where the continuous red line represents the speed limit $v^{MAX}(s)$.

Nodes in regions R_1 and R_4 of this graph (i.e., directly exceeding the limits) have been excluded from $\mathscr{N}$ already due to our definition of $\mathscr{V}_i$. Moreover, when the speed limit increases from 70 to 160 km/h at $s = 1$ km, the train would need some time and space to reach a speed matching the next higher limit. Therefore, we determine the maximal acceleration profile between these two speed values starting at $s = 1$ km, which corresponds to the dashed blue line in the figure. Nodes above this curve (i.e. in region R_2) are not reachable as they would imply an acceleration higher than a^{MAX} and are thus discarded from $\mathscr{N}$. Similarly, the dotted green line represents the maximal deceleration profile from 160 to 50 km/h computed so that the train reaches 50 km/h exactly when such limit is enforced at $s = 4$ km. Nodes in R_3 can be deleted from $\mathscr{N}$ because the train cannot reduce speed fast enough to meet the next speed limit. In conclusion we construct, for the locations on the track with a speed limit variation, a set of partial speed profiles using a^{MAX} or a^{MIN} as in Fig. 6,

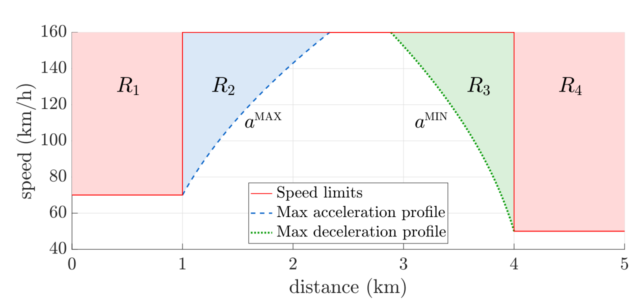

Fig. 6. Unreachable distance-speed regions due to speed limits.

and delete points that lie above such profile from $\mathscr{N}$. This enables further reducing the size of our distance-speed network $\mathscr{G}$ without losing feasible train trajectories (about 4% fewer nodes and arcs in the network used in our numerical study).

3.1.2. Arc set

We construct the arc set $\mathscr{A}$ as follows. For each pair of nodes n_1, $n_2 \in \mathscr{N}$, we add the arc $b = (n_1, n_2)$ to $\mathscr{A}$ if and only if: (i) n_1 and n_2 correspond to consecutive locations, respectively s_i and s_{i+1}, and (ii) the latter node is reachable from the previous node when taking into account all physical constraints on train motion. These constraints include not only the maximum acceleration and deceleration, as shown in Fig. 5, but also the maximum traction force and engine power. We refer to Section 2.1 for more details on these constraints. In other words, an arc $b \in \mathscr{A}$ in our graph represents a feasible train movement (hence a control) between two distance-speed points in consecutive locations.

In the following, we assume a uniform train acceleration over individual arcs, which seems a reasonable assumption when the distance discretization is sufficiently fine. Suppose that arc $b \in \mathscr{A}$ connects node $n_1 = (s_i, v_j)$ with $n_2 = (s_{i+1}, v_h)$. From our assumption it follows that each arc $b \in \mathscr{A}$ can be associated with a unique travel time t_b [s], an acceleration a_b [m/s^2], a force f_b [N], and the energy $e_b = E_{s_i}^{s_{i+1}}$ [Wh] consumed while traversing the arc. We compute these quantities using the standard equations of the uniformly accelerated motion, as well as, a discrete version of the train Equation (1). Specifically, time and acceleration of arc b are calculated with

$$t_b = \frac{2(s_{i+1} - s_i)}{v_j + v_h}, \qquad a_b = \frac{v_j + v_h}{t_b}, \tag{9}$$

and are not affected by the train resistance since distance and speed are fixed points in the network. Instead, the force f_b and consequently the energy e_b needed for the train to change its speed between two target values in a given distance interval depends on the train resistance, hence on the wind parameters w and θ (see also our related discussion at the end of Section 2.1). We state this dependency explicitly by writing $f_b = f_b(w, \theta)$ and $e_b = e_b(w, \theta)$. In principle, the force is also a function of the location s (i.e., it should be $f_b(s, w, \theta)$) but we compute it as a constant quantity over arc b as explained next. The wind (w, θ) is assumed to be constant in each interval $[s_i, s_{i+1}]$. By replacing the uniform acceleration a_b into Equation (1), we have:

$$f_b(s, w, \theta) = \rho\ m\ a_b + R'_{train}(v(s), w, \theta) + R_{line}(s), \quad \forall s \in [s_i, s_{i+1}], \tag{10a}$$

$$\simeq \rho\ m\ a_b + R'_{train}\left(\frac{v_j + v_h}{2}, w, \theta\right) + R_{line}\left(\frac{s_i + s_{i+1}}{2}\right), \quad \forall s \in [s_i, s_{i+1}], \tag{10b}$$

$$=: f_b(w, \theta), \quad \forall s \in [s_i, s_{i+1}]. \tag{10c}$$

Consequently, by applying (3) we obtain:

$$e_b(w, \theta) = \int_{s_i}^{s_{i+1}} \max\{f_b(w, \theta),\ 0\}\ ds \tag{11a}$$

$$= (s_{i+1} - s_i) \cdot \max\{f_b(w, \theta),\ 0\}. \tag{11b}$$

In (10b), the train and line resistances are determined, respectively, using the average speed over the interval and the line information at the midpoint $(s_i + s_{i+1})/2$. On one hand, in this equation we are introducing an approximation because these resistances in reality change continuously within the interval $[s_i, s_{i+1}]$ due to the continuous change of speed and possibly gradient and curvature of the track. (We expect the approximation error to be small when using sufficiently small distance intervals.) On the other hand, this strategy avoids the numerical approximation of integrals which are potentially time consuming (these integrals arise when combining (1a) with (3) and integrating over s). In fact, as we see in the next section, the force and energy consumption in (10)-(11) have to be computed for the entire network many times in order to obtain the desired energy-efficient trajectory.

3.2. Dynamic programming-based line search

Let us consider a wind scenario in which speed w and direction θ are constant in location and time, that is, wind conditions do not change during the journey that we aim to optimize (we explain later how to generalize this case). Our goal is to determine the speed profile that minimizes energy consumption while respecting the arrival time and accounting for (w, θ). Enforcing the arrival time means ensuring the trip takes less than a given amount of time that we call T^S [s]. The main challenge in tackling this problem is to handle jointly energy consumption and travel time, which can be seen as two conflicting objectives. Indeed, an early arrival time is usually associated with a higher average train speed, hence a higher energy consumption. Vice versa, if the train runs at low average speed, it will consume less energy but it will also take more time to reach destination.

Suppose for a moment that we want to minimize travel time alone without looking at energy. It is well known that finding the fastest train trajectory in our graph $\mathscr{G}$ is equivalent to solving a shortest path problem from node $n_{START} = (0, 0)$ to $n_{END} = (S, 0)$ (i.e., departure and arrival stations with zero speed) with arc costs equal to the travel time $c_b = t_b$ computed as in (9), for all $b \in \mathscr{A}$. This shortest path can be determined in $\mathscr{O}(|\mathscr{N}| + |\mathscr{A}|)$ using a simple forward DP (Ahuja et al., 1993), i.e., more efficiently than Dijkstra because $\mathscr{G}$ is acyclic. If the timetable is defined properly, a train following the fastest trajectory will take less than the scheduled time T^S to arrive at destination. However, this trajectory also leads to high energy consumption because the train is never coasting, for instance. Therefore, although it may be useful when the train has an initial delay, choosing the fastest trajectory is not desired during normal operations. On the other extreme of this trade-off, one could find the trajectory that minimizes energy consumption regardless of time. To do this, we assign arc costs equal to the energy needed to traverse the arc, i.e. $c_b(w, \theta) = e_b(w, \theta)$, following Equation (11), and solve a shortest path problem as a DP analogously to the previous case. Obviously this trajectory would be excessively slow (basically, driving at the minimum speed all time) and will result in a large delay. How can we then deal with both objectives? Notice that this challenge is not specifically wind-related but also arises when solving the standard train trajectory optimization problem.

One option to account for both objectives is to formulate our problem as a resource-constrained shortest path problem (RCSP; see, e.g., Hassin, 1992 or Reinhardt and Pisinger, 2011). The RCSP is a variant of the standard shortest path problem in which the consumption of an unconstrained resource has to be minimized while respecting an upper bound on the consumption of a second constrained resource. In our case, energy is the unconstrained resource and time the constrained resource. This problem is well known in the operations research literature and can be solved using a label-correcting algorithm (Irnich and Desaulniers, 2005). In brief, this algorithm stores in each iteration a set of *labels* for each node $n \in \mathscr{N}$, where a label is defined as a non-dominated (Pareto-optimal) time-energy pair corresponding to a path from the origin node n_{START} to the current node n. In each iteration, a node is selected and its labels extended to the successor nodes through a resource-extension function. Labels that violate the bound on the constrained resource are infeasible and hence discarded. The algorithm stops when no node can generate new non-dominated and feasible labels. Unfortunately, the RCSP is $\mathscr{NP}$-hard (Garey and Johnson, 1990) since the number of generated labels grows exponentially with the size of the network. In practice, solving an RCSP on large networks often requires excessive computational memory to store all the labels. To make the RCSP tractable, one could discretize a resource, e.g., time, into H steps and consider times falling within the same step to be equivalent. Under this assumption, the number of non-dominated labels for each node is bounded by H and the problem becomes weakly $\mathscr{N}\mathscr{P}$-hard, i.e., solvable in pseudo-polynomial time. However this approach introduces a potentially large approximation error when deleting dominated labels, with a magnitude that depends on H.

In our solution approach we would like, at the same time, to avoid the $\mathscr{N}\mathscr{P}$-hard complexity of the RCSP and also avoid introducing an additional discretization (since we already have discrete locations and speed values). Moreover, we would like to exploit the favorable $\mathscr{O}(|\mathscr{N}|+|\mathscr{A}|)$ complexity of the DP algorithm for solving standard shortest paths on acyclic graphs. We account for these features by proposing an approach that is new in train trajectory optimization and combines DP with a line search framework. The procedure is outlined in Algorithm 1. The core idea is pretty simple and consists in iteratively solving shortest paths on a network with arc costs defined as weighted sums:

$$c_b(\eta, w, \theta) = t_b + \eta \cdot e_b(w, \theta),\ \forall\, b \in \mathscr{A}, \tag{12}$$

i.e., linear combinations of travel time and energy consumption governed by a trade-off parameter $\eta \geq 0$. If $\eta = 0$, then $c_b(\eta, w, \theta) = t_b$ and the problem reduces to finding the fastest train trajectory. As $\eta > 0$ increases, more and more emphasis is put on energy consumption until the train will run late. The goal is hence optimizing the value of η such that the trajectory resulting from the DP matches as close as possible the target time T^S. We find such η value using a line search scheme. Specifically, at each iteration, we consider the current value for η, update all costs in the graph according to (12), and solve the corresponding DP. If the resulting trajectory takes a time T_η that is less than T^S, then we increase η to put more emphasis on energy; if it takes more time than T^S, then we decrease η to put more emphasis on travel time. The comparison between T_η and T^S thus guides the direction of the search by increasing or decreasing the value of η dynamically. After each iteration, the interval in which the desired η value lies is halved. The algorithm stops when a maximum number of iterations is reached or when a trajectory that matches T^S within a given time tolerance ε is found.

Algorithm 1. Line search DP for train trajectory optimization.

Inputs: Graph $\mathcal{G}$; Wind (w, θ); $\eta^{\text{MAX}} > 0$ (high value); Maximum iterations I; Arrival time tolerance ϵ.
Initialization: $T(w,\theta) = +\infty$, $E(w,\theta) = +\infty$, $\eta^{\text{MIN}} = 0$.
For iteration $i = 1$ to I **do**:

1. Set $\eta := (\eta^{\text{MAX}} + \eta^{\text{MIN}})/2$ and $c_b(\eta, w, \theta) := t_b + \eta\, e_b(w, \theta)$, $\forall b \in \mathcal{A}$, obtaining graph $\mathcal{G} = \mathcal{G}(\eta, w, \theta)$;
2. Solve shortest path as a DP on $\mathcal{G}(\eta, w, \theta)$, resulting in trajectory X_η, travel time T_η, and energy E_η;
3. **If** $|T_\eta - T^{\text{S}}| < |T(w,\theta) - T^{\text{S}}|$, update current best solution $X(w,\theta) = X_\eta$, $T(w,\theta) = T_\eta$, $E(w,\theta) = E_\eta$;
4. **If** $T_\eta < T^{\text{S}}$, redefine $\eta^{\text{MIN}} = \eta$, **else**, redefine $\eta^{\text{MAX}} = \eta$;
5. **If** $|T(w,\theta) - T^{\text{S}}| < \epsilon$, **break**.

Outputs: Optimized train trajectory $X(w,\theta)$, time $T(w,\theta)$, and energy consumption $E(w,\theta)$ for wind scenario (w,θ).

From a multi-objective optimization theory perspective, the mentioned RCSP approach and our proposed line search method exhibit different properties. The RCSP belongs to the class of *ε-constrained* methods in which only one of the objectives is kept for minimization and the others are turned into constraints. While ε-constrained methods are able to generate all Pareto-optimal (i.e., non-dominated) solutions, they typically add complexity to the problem (Ehrgott, 2006). This is also the case in our application since the RCSP is $\mathscr{N}\mathscr{P}$-hard as discussed earlier. In contrast, our line search method uses a trade-off parameter to combine the two objectives into a single objective, which is a scalarization technique commonly referred to as the *weighted sum* method. Although this technique is guaranteed to produce only Pareto-optimal solutions, it also has a well-known drawback: it can only find nondominated points on the convex hull of the nondominated set, meaning that if the Pareto frontier contains non-convexities, then there exist nondominated solutions that are unreachable by using a weighted sum of the objectives as a single objective (Boland et al., 2015; Ehrgott, 2006). In our numerical study, we show that this drawback is not problematic in our application by approximating the Pareto frontier and noticing that there are no big gaps between consecutive points that would indicate the possibility non-convexities, hence unreachable points. Therefore, we can say that our algorithmic choice is not only supported by its computational efficiency but also by its capability to approximate the Pareto frontier of our bi-objective problem.

The described line search method can be extended to handle location-dependent wind information or forecasts in a straightforward manner. Assume that wind speed and direction are given as vectors $w \in \mathbb{R}^I_+$ and $\theta \in [0, 2\pi]^I$, respectively, where $(w(i), \theta(i))$ is the wind in interval $[s_i, s_{i+1}]$. For example, this vector could represent the current wind conditions measured along the track and supposed not to change in the short term. Alternatively, this vector could represent a forecast vector, where each component $(w(i), \theta(i))$ is the wind forecast for the interval $[s_i, s_{i+1}]$ when the train will cross that location, computed using the current wind information. Either way, the wind data (i.e., measurements or forecasts) used in the algorithm is known in advance and is static in time. Handling dynamic changes in wind and the related flexibility to adjust the trajectory in real time is not the purpose of this paper. For each arc $b \in \mathscr{A}$, we compute travel time and acceleration as before using Equation (9). Then, we compute the force $f_b(w(i), \theta(i))$ and consequent energy $e_b(w(i), \theta(i))$ following Equations (10) and (11) but using the specific wind information $(w(i), \theta(i))$ available for that arc. At this point, Algorithm 1 can be applied directly to the network.

We can evaluate the quality of a single wind-aware trajectory for a given wind scenario (w, θ) by following these four steps:

1. Solve Algorithm 1 for scenario (w, θ), resulting in our wind-aware trajectory $X(w, \theta)$ with energy consumption $E(w, \theta)$ [kWh];
2. Solve Algorithm 1 by setting $(w, \theta) = (0, 0)$, resulting in the traditional "no-wind" trajectory $X(0, 0)$ with energy consumption $E(0, 0)$ [kWh];
3. Evaluate the no-wind trajectory $X(0, 0)$ on the actual wind condition (w, θ), i.e., using the true arc energy consumption $e_b(w, \theta)$, for all $b \in \mathscr{A}$. This results in a different energy consumption for this trajectory, denoted $\tilde{E}(w, \theta)$ [kWh];
4. Compare $E(w, \theta)$ with $\tilde{E}(w, \theta)$.

It is important to compare $E(w, \theta)$ with $\tilde{E}(w, \theta)$ and not with $E(0, 0)$, that is, the two trajectories should be evaluated on the same wind conditions. Moreover, it is important to use a low tolerance ε on arrival time so that the difference between the two consumption levels can be fully attributed to the knowledge and optimization of wind information. To further assess the quality of our trajectories in a statistically sound manner, one could generate a set of K wind scenarios $\{(w_k, \theta_k), k \in \mathscr{K}\}$ and, for each $k \in \mathscr{K}$, follow the steps mentioned above. This procedure results in two energy consumption distributions $\{E(w_k, \theta_k),\ k \in \mathscr{K}\}$ and $\{\tilde{E}(w_k, \theta_k),\ k \in \mathscr{K}\}$, that we will compare in Section 4.2.

4. Numerical study

In this section, we present our numerical experiments and findings. We describe the instance and computational setting in Section 4.1. We compute energy-efficient trajectories under different wind scenarios in Section 4.2. Finally, we perform a sensitivity analysis on some of the most relevant problem parameters in Section 4.3.

4.1. Instance and computational setup

In Table 1, we report the train characteristics that we use in our experiments. Almost all of these parameters are taken from Wang and Goverde (2016a) and refer to a local train from the Dutch railway: a Sprinter. Since no empirical data is available to estimate the two scaling

Table 1
Train and resistance parameters.

Name	Value	Unit	Name	Value	Unit
m	220	t	p^{MAX}	1918	kW
ρ	1.06	–	f^{MAX}	170	kN
L	138	m	c_1	5.8	kN
l	3.2	m	c_2	0.072	kN s/m
h	4.4	m	c_3	0.013	kN $(s/m)^2$
a^{MAX}	0.6	m/s^2	ξ_1	1.0	–
a^{MIN}	−0.8%	m/s^2	ξ_2	0.01	–

parameters ξ_1 and ξ_2, we select $\xi_1 = 1$ so that our resistance model is consistent with the model from the literature (5) without cross-wind, and a small value for $\xi_2 = 0.01$. We will discuss in our sensitivity analysis in Section 4.3 how the results change when varying the value of these parameters.

We use a 20 km long synthetic track containing six different speed limits displayed in Fig. 7. Line resistances (e.g., track slopes and curves) can be handled by our method as discussed in Section 3.1.2 but we set them to zero to better decouple the effect of wind on trajectories when interpreting the results. This figure shows as well the discretization parameters that we choose. In particular, we use 201 uniformly distributed locations so that the control intervals are equal to $\Delta s = 100$ m, and 300 speed values generated according to Equation (8) with $\varphi = 0.6$. The resulting graph $\mathscr{G}$ has approximately 46 000 nodes and 2.8 million feasible arcs. When running Algorithm 1, we choose $\eta^{MAX} = 10$ because we know from experiments that this value results in trajectories with travel time longer than T^S. In other words, is always larger than the optimal trade-off value η in Equation (12) resulting from Algorithm 1 (usually between 3 and 3.5 in our tests), which is a condition needed for the line search procedure to be executed. Then, we set $I = 100$ and $\varepsilon = 0.5$, that is, our trajectories are within 0.5 s from the target T^S. This is essential because, as mentioned, it makes sense to compare energy consumption only for trajectories that exhibit the same (or very similar) travel time. Otherwise, a different energy use could be a mere consequence of the fact that slower trains consume less energy on average.

The algorithm was implemented in C++ and run on a standard laptop equipped with a i7-8650U processor with 16 GB RAM. For a given wind scenario (w, θ), the running time needed to compute one train trajectory, i.e., executing the full line search algorithm, was 5.5 s on average. More specifically, solving one DP (single η value) took on average 0.4 s. This time includes 0.32 s of preprocessing time to update all arc costs in the network given η, and 0.08 s to run the DP itself. The line search algorithm converged on average in 13 iterations.

4.2. Results

We start by investigating whether our line search method is able to approximate in a satisfactorily manner the Pareto frontier of our bi-objective optimization problem. We know in fact from Section 3.2 that our method produces solutions which are guaranteed to be Pareto-optimal, but it might not be able to obtain all Pareto-optimal solutions. To investigate this potential shortcoming, we approximate three Pareto frontiers corresponding to three different wind scenarios: (i) $w = 0$, i.e., the standard trajectory optimization problem without wind, (ii) $w = 7$ m/s and $\theta = 0$, and (iii) $w = 7$ m/s and $\theta = 3/4\pi$. We assume the wind is constant throughout the duration of travel, i.e. $(w(i), \theta(i)) = (w, \theta)$ for all $i = 0, \ldots, I-1$. To approximate a frontier, we varied the scheduled trip time T^S from 840 to 1040 s with steps of 1 s, and run Algorithm 1 for each of these T^S values obtaining a collection of trajectories with given travel time and energy consumption. The resulting frontiers reported in Fig. 8 present no large time/energy gaps between consecutive points. The maximum gap between time values is 3.1 s, which decreases to less than 2 s when restricting to the interval $T^S \in [870, 920]$ around our reference value of $T^S = 900$. This analysis suggests that our line search approach is particularly suitable to efficiently generate and approximate the Pareto frontier in the train trajectory optimization problem, without requiring more sophisticated bi-objective optimization techniques.

The frontiers in Fig. 8 also highlight the importance of exploiting the available running time for the purpose of saving energy. For example, on average across the three frontiers, optimal train trajectories based on $T^S = 900$ seconds reduce consumption by 6.8% (7.4 kWh) and 14.2% (16.7 kWh), respectively, compared to the faster trajectories computed for 880 and 860 s.

Next, we examine the impact of wind by determining the wind-aware trajectories for different combinations of wind speed and direction. Specifically, we consider 30 values of w ranging from 0 m/s to 15 m/s, and ten values of θ in the interval $[0, \pi]$. We do not consider values of θ larger than π because the effect of wind is symmetric with respect to the direction of the track, i.e., the trajectories computed for $(w, \pi + \theta)$ and $(w, \pi - \theta)$ would be equivalent. In Fig. 9, we show for each (w, θ)-pair the energy saving achieved when accounting for wind information. Precisely, each value in this graph corresponds to the percentage difference between the energy use in the classical no-wind trajectory $\tilde{E}(w, \theta)$ and the energy use in our wind-aware trajectories $E(w, \theta)$. We denote this energy saving [%] by $\mathscr{E}(w, \theta) := [\tilde{E}(w, \theta) - E(w, \theta)]/\tilde{E}(w, \theta) \cdot 100$.

From the figure, it is noteworthy observing that the energy saving $\mathscr{E}(w, \theta)$ increases with the wind speed w. This behavior was somewhat expected: the strongest the wind, the more relevant accounting for wind information would be. For example, the average energy saving across wind directions is 1.6% and 4.5%, respectively, for wind speeds of 7 m/s and 14 m/s. The maximum energy saving achieved can be up to 10–11% under strong winds. We believe that these saving figures are remarkable considering that we are competing with trajectories which are already energy-efficient and state-of-the-art, with the only difference of exploiting wind information available before departure. Furthermore, the energy saving $\mathscr{E}(w, \theta)$ varies substantially with the wind direction. In particular, the highest savings are achieved under headwinds or winds hitting from a direction of $\theta \in [0, \pi/6]$, followed by $\theta = \pi/4$. For

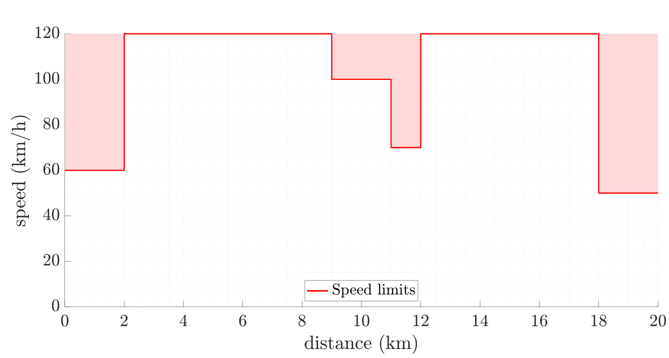

Name	Value	Unit
S	20	km
Δs	100	m
$\lvert\mathcal{S}\rvert$	201	-
$\lvert\mathcal{V}\rvert$	300	-
V^{MAX}	120	km/h
T^S	900	s
ϕ	0.6	-

Fig. 7. Illustration of the track and discretization parameters.

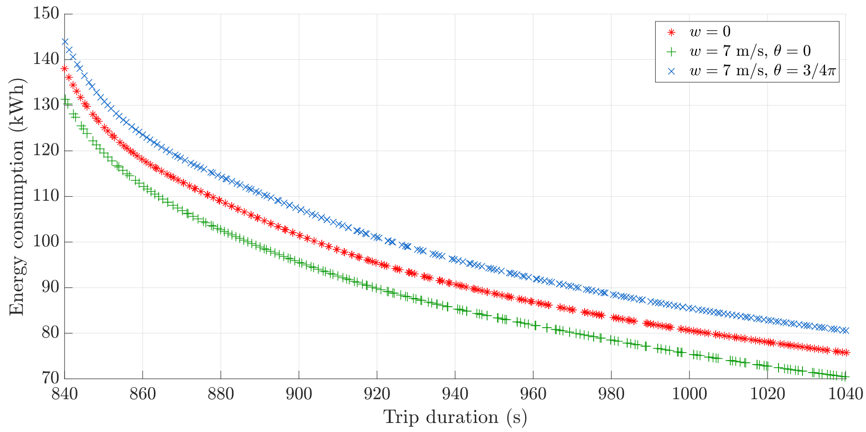

Fig. 8. Approximations of travel time and energy consumption Pareto frontiers.

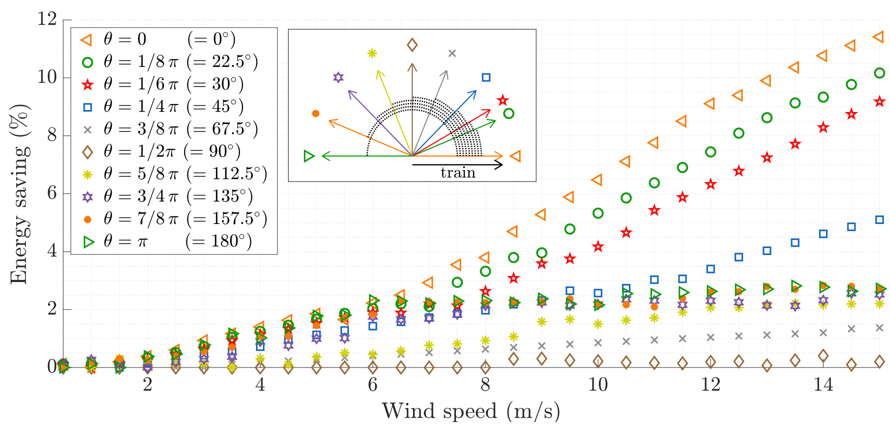

Fig. 9. Energy savings achieved with different wind speeds and directions.

$\theta \in [3/8\pi, \pi]$, the saving is smaller and does not exceed 3% even when the wind is strong. This hints a relationship between the wind direction and the flexibility to manipulate the trajectories to reduce energy consumption. This relationship is however complex to disentangle as it involves the equations of relative train speed and yaw angle (4), wind-dependent drag coefficient (6), and resistance (7), all combined with the optimization process.

Winds of 12–15 m/s as in Fig. 9 are rather strong and might be uncommon in several regions. In most European countries, for instance, average wind speeds range from 4 to 7 m/s (Global Wind Atlas, 2019). Therefore, to obtain statistics on potential energy savings, we evaluate the no-wind and wind-aware trajectories using $K = 500$ wind scenarios generated as follows. Wind speed is modeled with a Weibull distribution, which is a common choice in the wind energy literature (Hennessey Jr., 1977). We choose shape and scale parameter of the Weibull equal to 2 and 6, respectively, resulting in the wind speed distribution displayed in Fig. 10(a). Wind direction is sampled from a uniform distribution $\theta \tilde{} U[0, 2\pi]$. For each scenario k, we compute the corresponding

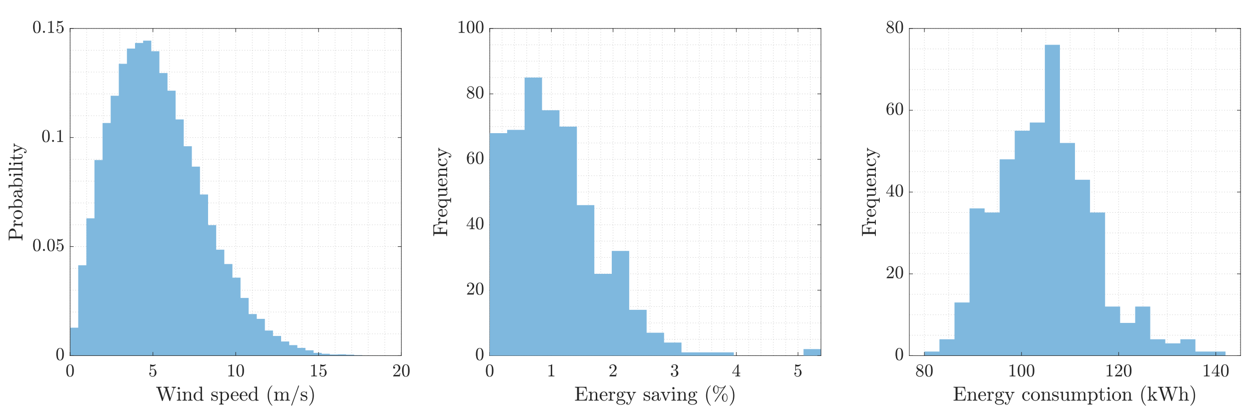

(a) Wind speed distribution. (b) Energy consumption histogram. (c) Energy saving histogram.

Fig. 10. 500 wind speed scenarios and resulting energy consumption energy saving.

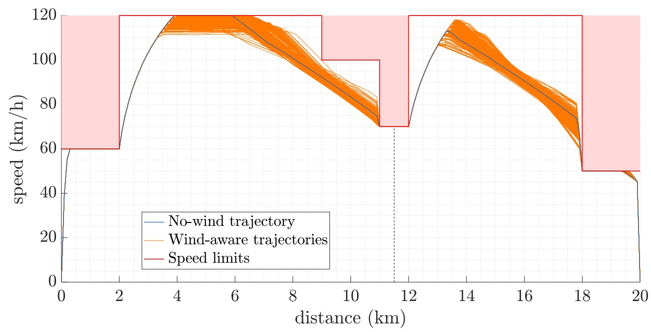

Fig. 11. Wind-aware train trajectories computed for the 500 wind scenarios of Fig. 10(a).

wind-aware trajectories and report in Fig. 10(b) and (c) the histograms, respectively, for $\mathscr{E}(w,\theta)$, i.e. the energy saving from wind-aware trajectories, and their energy consumption.

The average energy saving from the distribution in Fig. 10(b) is 1.1% (corresponding to 1.2 kWh), and 95% of the scenarios lead to a saving below 2.5%. This is consistent with the results previously shown in the left part of Fig. 9 because most scenarios have a wind speed lower then 8 m/s (see Fig. 10(a)). Even though these saving levels might seem a small amount of energy, summing them up over all train trips on a network during a day may result in a significant amount. The distribution shown in Fig. 10(c) reveals that the energy consumption of a train is affected substantially by the wind conditions. While the average energy consumption is 105 kWh, the minimum and maximum across the 500 scenarios, respectively, are equal to 81 kWh and 140 kWh. This observation further underscores the important role of wind on the energy consumption of a train run.

In Fig. 11 we show the wind-aware trajectories corresponding to our 500 scenarios. The wind-aware trajectories deviate from the original trajectory and, as expected, the deviations mainly occur for high train speeds for which the air resistance is higher, and consequently the potential impact of wind is higher. Furthermore, we notice that the new trajectories mainly differ in the shape of the coasting profile and the moment in which the coasting phases start, whereas the acceleration profile is never altered. The reason for the acceleration profile not changing is that, regardless of wind, it is always preferable to use maximum acceleration and subsequently exploit coasting (where consumption is zero) rather then keep accelerating the train at a slower rate. This follows from the classical analytical results on train trajectory optimization based on the Pontryagin's principle (Cheng and Howlett, 1992; Howlett, 2000). The wind affects the resistance and hence the optimal switching points, e.g., between maximum acceleration and coasting, and also the shape of the coasting profile, but the aforementioned property related to maximum acceleration holds regardless of wind.

Next, we show how our method handles location-dependent wind information, i.e., the wind speed and direction vary in different track locations. A key reason for wind to vary with location is that the direction of the track changes. For example, if the wind is from the north and the track heads north and then north-east, the apparent wind direction will change. Location-dependent wind data can be interpreted as wind measurements along the track, or a forecast vector computed at the beginning of the journey and describing the expected wind (see also the related discussion in Section 3.2). To derive some qualitative insights on the trajectories, we focus on a simple and illustrative example with two distinct wind zones: the first part of the track $[0,\widehat{S}]$ presents some wind conditions (w^A,θ^A) while the second part $(\widehat{S},S]$ some different conditions (w^B,θ^B). Specifically, we assume that a strong wind of 10 m/s blows according to the two scenarios shown in Fig. 12. We compute the wind-aware trajectory for both scenarios and compare them with the no-wind trajectory in Fig. 13.

In the first scenario in Fig. 12(a), the wind is initially opposed to the train motion (i.e., resistance increases) but afterwards it becomes in favor (i.e., resistance decreases). From Fig. 13, we can see that the corresponding wind-aware trajectory (dashed green line) deviates from the no-wind trajectory (continuous blue time), especially for high train speed values. For example, in the first part of the trip, due to higher resistances it is optimal for the train: (i) not to reach the speed limit but to hold instead a slightly lower speed, and (ii) to start coasting later. In the second part of the trip, the shape of the coasting profile changes since the wind is favorable, hence the train speed during coasting decreases at a slower rate. In the second scenario in Fig. 12(b), the wind conditions are precisely the opposite and, again, the wind-aware profile (dotted orange line) deviates from the no-wind one. The same observations made for the first scenario holds for the second scenario too on opposite zones. The energy saving is 3.8% on average across the two scenarios.

4.3. Sensitivity analysis

Although most parameters of our baseline instance in Table 1 refer to a real railway vehicle, the scaling values ξ_1 and ξ_2 were introduced with our model of train resistance (6)-(7). As discussed at the end of Section 2.2, one would expect that ξ_1 takes a value close to 1 for which, in the absence of cross-wind, our resistance equation (7) reduces to the simpler equation from the literature (5). Moreover, one would expect that ξ_2

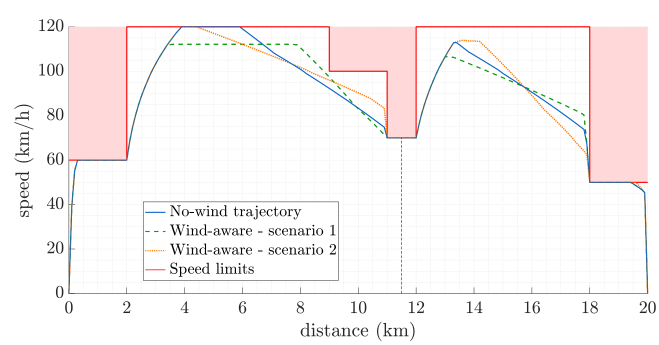

Fig. 13. Wind-aware vs. no-wind train trajectories under the scenarios in Fig. 12.

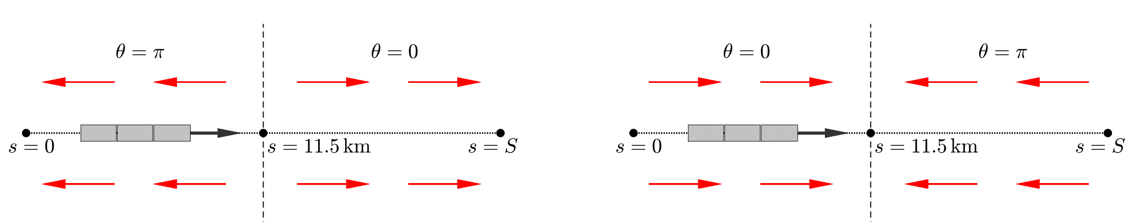

(a) Scenario 1 (b) Scenario 2

Fig. 12. Scenarios with location-dependent wind information (two wind zones).

takes a much smaller value than ξ_1 since the effect of the former parameter on (6) scales with the train lateral surface, which is substantially larger than the frontal surface. In this analysis, we assess the energy consumption and saving of wind-aware trajectories for different choices of ξ_1 and ξ_2. Based on experiments, we consider 10 equidistant values for ξ_1 and ξ_2 ranging, respectively, from 0.92 to 1.10 and from 0 to 0.027. For each combination of them (i.e., a total of 100 combinations), we compute the no-wind and wind-aware trajectory for the case of constant wind of $w = 10$ m/s and $\theta = \pi/4$. We display the resulting energy consumption values $E(w,\theta)$ [kWh] and energy saving values $\mathscr{E}(w,\theta)$ [%] in the form of heatmaps in Fig. 14

As expected, Fig. 14(a) shows that the energy consumption increases approximately linearly when either ξ_1 or ξ_2 is increased. When ξ_2 is fixed, increasing ξ_1 from 0.92 to 1.10 results in an increase in energy consumption of 3.8 kWh on average (i.e., about 4%). Vice versa, when ξ_1 is fixed and ξ_2 is increased from 0 to 0.027, the average consumption increases on average by 8.5 kWh (i.e., almost 10%), which is significant. The energy saving in Fig. 14(b) also depends on ξ_1 and ξ_2 and is generally larger when these parameters are set to higher values. This behavior is explained by the fact that higher ξ_1 and ξ_2 values imply a stronger impact of wind on train resistance, hence knowing wind information and optimizing trajectories accordingly becomes more valuable. Overall, our analysis suggests that the impact of wind on energy consumption and saving of wind-aware trajectories can be significant, but future empirical research estimating ξ_1 and ξ_2 from data might help obtaining more accurate estimates of energy consumption and saving.

We also tested the robustness of our method and trajectories towards different discretization choices. We considered 10 values for the number of discrete locations $|\mathscr{S}|$ ranging from 100 to 280, and 10 values for the number of speed points $|\mathscr{V}|$ from 125 to 350. This gives a total 100 different networks that we tested for the case of constant wind of $w = 10$ m/s and $\theta = \pi/4$. From the results, we can conclude that our method is rather robust towards changes in the underlying network. In fact:

1. The wind-aware trajectory remains quite stable and the associated energy consumption only varies between 87.0 and 90.2 kWh, with variance 0.67 kWh.
2. The running time to execute the line search algorithm varies between 1.2 and 16.5 s. Therefore, the algorithm remains tractable and fast when increasing the number of locations and speed points. Notice that, for a fixed number of speed points $|\mathscr{V}|$, increasing the number of locations $|\mathscr{S}|$ does not necessarily result in larger networks. Indeed, for a given node at location s_i, the number of reachable nodes at s_{i+1} (hence feasible arcs in $\mathscr{G}$) decreases when the distance $s_{i+1} - s_i$ decreases, as can be seen from Fig. 5 by taking vertical sections of the graph.

5. Conclusion

In this paper, we have developed a train speed optimization model accounting for external uncontrollable variables that manifest differently across journeys on the same track, such as weather conditions. In particular, we considered wind variability and proposed a new train resistance equation that is grounded on the literature on aerodynamic drag but accounts for wind speed and direction more explicitly by decoupling the effect of wind on the traversal and longitudinal surfaces of the train. Using this equation, we optimized the train speed profile under a given wind scenario by combining line search with a dynamic programming-based shortest path algorithm defined on a distance-speed network. Our network is based on a novel non-linear discretization of speed values that makes it more balanced, and our solution approach is also new in train trajectory optimization and computationally efficient. We performed a numerical study where we tested the algorithm under different wind conditions. The results revealed that our wind-aware trajectories outperform classical trajectories computed without using wind data before train departure but only the available running time from the timetable. In particular, the energy saving depends on the wind direction and is usually about 0–2% for low wind speed values but can reach up to 10% for strong headwinds.

This paper is the first showing the benefit of designing static proactive train speed profiles incorporating external factors that depend on the specific journey conditions. If combined with a DAS, railway operators may benefit from our model by using it: (i) on a tactical level, by redefining the trajectory given an average wind profile, or (ii) in near real time, i.e., right before the trip starts, to update the trajectory based on the current or forecast weather, especially under strong winds.

This paper opens multiple avenues for future research. First, an empirical validation of the proposed resistance equation and the calibration using real data of the scaling parameters in this equation (ξ_1 and ξ_2) would help obtaining more accurate estimates of the energy saving. Second, it could be relevant to incorporate more external, journey-dependent factors in the model such as rain conditions and train load. Indeed we expect the energy saving to be larger when computing profiles using more factors/information. Finally, despite handling location-dependent wind data proactively, our trajectories are static in time and computed at the beginning of the trip assuming that the wind measurements or forecasts will not change in the short term. Therefore, an extension of this work could be modeling wind as a stochastic process that evolves dynamically over time, allowing the trajectory to adjust in

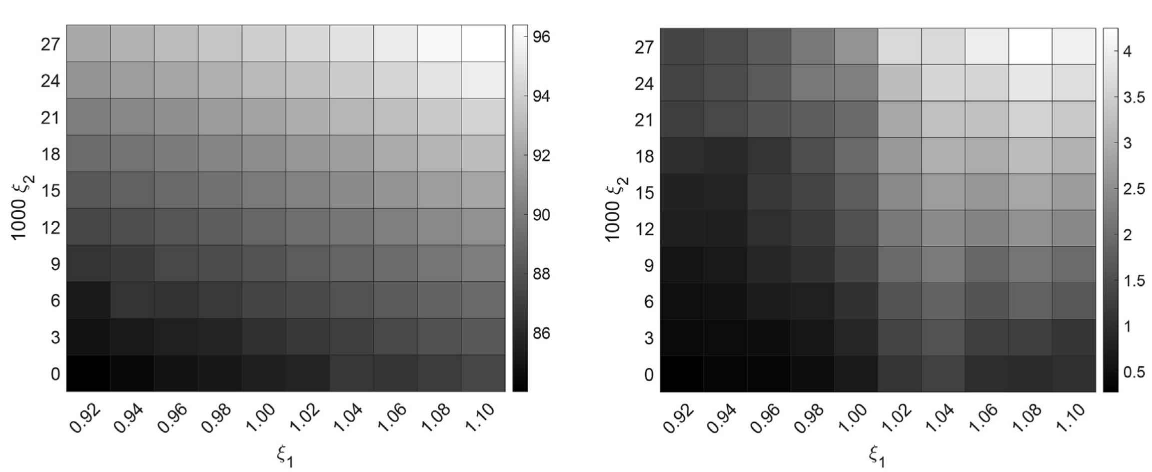

(a) Heatmap of energy consumption (b) Heatmap of energy saving

Fig. 14. Heatmap of energy consumption and saving under different ξ_1 and ξ_2 combinations.

real time. This extension is perhaps more relevant for longer train trips, such as in high speed railways, in which wind conditions might change more.

Declaration of competing interest

The Authors have no interests to declare.

Acknowledgments

Alessio Trivella and Francesco Corman are supported by the Swiss National Science Foundation under the research project DADA, grant no. 181210. Pengling Wang is supported by the NWO (Dutch Research Council) under the Rubicon Grant no. 019.172EN.010.

Appendix

Here we illustrate the steps to derive the relative train speed v_A and yaw angle γ in Equation (4) from the train speed v, real wind speed w, and wind impact angle θ. We consider the triangle with sides v, v_A and w as shown in Fig. 15. This figure also shows y, which is the projection onto v of the vertex made by sides v_A and w, and the two segments v_1 and v_2 generated by this projection.

We know that v_A can be computed from v_1 and y by $v_A = \sqrt{v_1^2 + y^2}$. Furthermore, it is easy to see that $y = w\sin\theta$, $v_2 = w\cos\theta$, and $v_1 = v - v_2$. It follows that $v_1^2 = (v - w\cos\theta)^2$. Replacing the expressions for y and v_1 into the one for v_A gives Equation (4a). Finally, the law of sines implies that $(\sin\gamma)/w = (\sin\theta)/v_A$, which enables expressing γ as function of the other quantities. This expression can be reformulated as in Equation (4b).

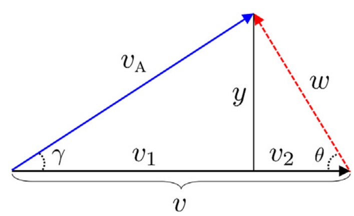

Fig. 15. Derivation of v_A and γ.

References

Ahuja, R., Magnanti, T., O, J.B., 1993. Network Flows: Theory, Algorithms, and Applications. Prentice Hall.
Albrecht, A., Howlett, P., Pudney, P., Vu, X., Zhou, P., 2016a. The key principles of optimal train control part 1: formulation of the model, strategies of optimal type, evolutionary lines, location of optimal switching points. Transp. Res. Part B Methodol. 94, 482–508.
Albrecht, A., Howlett, P., Pudney, P., Vu, X., Zhou, P., 2016b. The key principles of optimal train control part 2: existence of an optimal strategy, the local energy minimization principle, uniqueness, computational techniques. Transp. Res. Part B Methodol. 94, 509–538.
Boland, N., Charkhgard, H., Savelsbergh, M., 2015. A criterion space search algorithm for biobjective integer programming: the balanced box method. Inf. J. Comput. 27 (4), 735–754.
Bomhauer-Beins, A., 2019. Energy Saving Potentials in Railway Operations under Systemic Perspectives. Ph.D. thesis. ETH Zurich.
Cheng, J., Howlett, P., 1992. Application of critical velocities to the minimisation of fuel consumption in the control of trains. Automatica 28 (1), 165–169.
De Martinis, V., Corman, F., 2018. Data-driven perspectives for energy efficient operations in railway systems: current practices and future opportunities. Transport. Res. C Emerg. Technol. 95, 679–697.
Ehrgott, M., 2006. A discussion of scalarization techniques for multiple objective integer programming. Ann. Oper. Res. 147 (1), 343–360.
Franke, R., Terwiesch, P., Meyer, M., 2000. An algorithm for the optimal control of the driving of trains. In: Proceedings of the 39th IEEE Conference on Decision and Control, vol. 3, pp. 2123–2128.
Fs, 2019. Ferrovie dello Stato Italiane - Environmental Commitment. https://www.fsitaliane.it/content/fsitaliane/en/sustainability/environmental-commitment/energy.html. (Accessed 12 March 2020).
Garey, M., Johnson, D., 1990. Computers and Intractability; A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., New York, NY, USA.
Ghaviha, N., Bohlin, M., Holmberg, C., Dahlquist, E., Skoglund, R., Jonasson, D., 2017. A driver advisory system with dynamic losses for passenger electric multiple units. Transport. Res. C Emerg. Technol. 85, 111–130.
Global Wind Atlas, 2019. Global wind Atlas - wind speed. https://www.globalwindatlas.info/. (Accessed 12 March 2020).
González-Gil, A., Palacin, R., Batty, P., Powell, J., 2014. A systems approach to reduce urban rail energy consumption. Energy Convers. Manag. 80, 509–524.
González-Gil, A., Palacin, R., Batty, P., 2015. Optimal energy management of urban rail systems: key performance indicators. Energy Convers. Manag. 90, 282–291.
Gunselmann, W., 2005. Technologies for increased energy efficiency in railway systems. In: 2005 European Conference on Power Electronics and Applications. IEEE, p. 10.
Haahr, J.T., Pisinger, D., Sabbaghian, M., 2017. A dynamic programming approach for optimizing train speed profiles with speed restrictions and passage points. Transport. Res. C Emerg. Technol. 99, 167–182.
Hansen, I.A., Pachl, J., 2014. Railway Timetabling & Operations. Eurailpress, Hamburg.
Hassin, R., 1992. Approximation schemes for the restricted shortest path problem. Math. Oper. Res. 17 (1), 36–42.
Hennessey Jr., J.P., 1977. Some aspects of wind power statistics. J. Appl. Meteorol. 16 (2), 119–128.
Howlett, P., 2000. The optimal control of a train. Ann. Oper. Res. 98 (1-4), 65–87.
Howlett, P.G., Pudney, P.J., 2012. Energy-Efficient Train Control. Springer Science & Business Media.
IEA, 2019. International Energy Agency - Tracking Clean Energy Progress: Transport. https://www.iea.org/tcep/transport/. (Accessed 12 March 2020).
Irnich, S., Desaulniers, G., 2005. Shortest path problems with resource constraints. In: Column Generation. Springer, pp. 33–65.
Ko, H., Koseki, T., Miyatake, M., 2004. Application of dynamic programming to the optimization of the running profile of a train. In: Computers in Railway Six. WIT Press, pp. 103–112.
Lai, Y., Barkan, C., Önal, H., 2008. Optimizing the aerodynamic efficiency of intermodal freight trains. Transport. Res. E Logist. Transport. Rev. 44 (5), 820–834.
Liu, R., Golovitcher, I., 2003. Energy-efficient operation of rail vehicles. Transport. Res. Pol. Pract. 37 (10), 917–932.
Lu, S., Hillmansen, S., Ho, T., Roberts, C., 2013. Single-train trajectory optimization. IEEE Trans. Intell. Transport. Syst. 14 (2), 743–750.
Lu, S., Weston, P., Hillmansen, S., Gooi, H.B., Roberts, C., 2014. Increasing the regenerative braking energy for railway vehicles. IEEE Trans. Intell. Transport. Syst. 15 (6), 2506–2515.
Luan, X., Wang, Y., De Schutter, B., Meng, L., Lodewijks, G., Corman, F., 2018a. Integration of real-time traffic management and train control for rail networks-part 1: optimization problems and solution approaches. Transp. Res. Part B Methodol. 115, 41–71.
Luan, X., Wang, Y., De Schutter, B., Meng, L., Lodewijks, G., Corman, F., 2018b. Integration of real-time traffic management and train control for rail networks-part 2: extensions towards energy-efficient train operations. Transp. Res. Part B Methodol. 115, 72–94.
NS, 2019. Nederlandse Spoorwegen - Sustainability at NS: Energy Efficiency. https://www.ns.nl/en/about-ns/sustainability/climate-neutral/energy-saving-measures.html. (Accessed 12 March 2020).
Pachl, J., 2002. Railway Operation and Control. VTD Rail Publishing, Mountlake Terrace, USA.
Panou, K., Tzieropoulos, P., Emery, D., 2013. Railway driver advice systems: evaluation of methods, tools and systems. Journal of Rail Transport Planning & Management 3 (4), 150–162.
Peters, J.-L., 1990. Bestimmung des aerodynamischen widerstandes des ICE/V im tunnel und auf freier strecke durch auslaufversuche. ETR. Eisenbahntechnische Rundschau 39 (9), 559–564.
Powell, J., Palacín, R., 2015. A comparison of modelled and real-life driving profiles for the simulation of railway vehicle operation. Transport. Plann. Technol. 38 (1), 78–93.

Railenergy, 2016. Innovative integrated energy efficiency solutions for railway rolling stock, rail infrastructure and train operation. http://www.railenergy.org/background.shtml. (Accessed 12 March 2020).
Reinhardt, L., Pisinger, D., 2011. Multi-objective and multi-constrained non-additive shortest path problems. Comput. Oper. Res. 38 (3), 605–616.
Rochard, B., Schmid, F., 2000. A review of methods to measure and calculate train resistances. Proc. Inst. Mech. Eng. - Part F J. Rail Rapid Transit 214 (4), 185–199.
SBB, 2019. Energy Efficiency. https://company.sbb.ch/en/sbb-as-business-partner/services-rus/energy/energiesparen.html. (Accessed 12 March 2020).
Scheepmaker, G.M., Goverde, R., Kroon, L., 2017. Review of energy-efficient train control and timetabling. Eur. J. Oper. Res. 257 (2), 355–376.
Schetz, J.A., 2001. Aerodynamics of high-speed trains. Annu. Rev. Fluid Mech. 33 (1), 371–414.
Tian, H.-Q., 2009. Formation mechanism of aerodynamic drag of high-speed train and some reduction measures. J. Cent. S. Univ. Technol. 16 (1), 166–171.
UIC, 2019. International union of railways - energy and CO2 emissions. https://uic.org/sustainable-development/energy-and-co2-emissions/. (Accessed 12 March 2020).
Vardy, A., Reinke, P., 1999. Estimation of train resistance coefficients in tunnels from measurements during routine operation. Proc. Inst. Mech. Eng. - Part F J. Rail Rapid Transit 213 (2), 71–87.
Wang, P., Goverde, R.M.P., 2016a. Multiple-phase train trajectory optimization with signalling and operational constraints. Transport. Res. C Emerg. Technol. 69, 255–275.
Wang, P., Goverde, R.M.P., 2016b. Two-train trajectory optimization with a green-wave policy. Transport. Res. Rec. 2546 (1), 112–120.
Wang, J., Rakha, H.A., 2018. Longitudinal train dynamics model for a rail transit simulation system. Transport. Res. C Emerg. Technol. 86, 111–123.
Wang, Y., Ning, B., Cao, F., De Schutter, B., Van den Boom, T., 2011. A survey on optimal trajectory planning for train operations. In: Proceedings of 2011 IEEE International Conference on Service Operations, Logistics and Informatics. IEEE, pp. 589–594.
Wang, Y., De Schutter, B., van den Boom, T.J., Ning, B., 2013. Optimal trajectory planning for trains–a pseudospectral method and a mixed integer linear programming approach. Transport. Res. C Emerg. Technol. 29, 97–114.
Wang, P., Trivella, A., Corman, F., 2020. Train trajectory optimization for improved on-time arrival under parametric uncertainty. Transport. Res. C Emerg. Technol. (Forthcoming).
Yang, X., Li, X., Ning, B., Tang, T., 2016. A survey on energy-efficient train operation for urban rail transit. IEEE Trans. Intell. Transport. Syst. 17 (1), 2–13.
Ye, H., Liu, R., 2017. Nonlinear programming methods based on closed-form expressions for optimal train control. Transport. Res. C Emerg. Technol. 82, 102–123.
Zhou, L., Tong, L.C., Chen, J., Tang, J., Zhou, X., 2017. Joint optimization of high-speed train timetables and speed profiles: a unified modeling approach using space-time-speed grid networks. Transp. Res. Part B Methodol. 97, 157–181.