

Predicting the Impact of Disruptions to Urban Rail Transit Systems

Xiaoyun Mo¹, Chu Cao¹, Mo Li¹, David Z.W. Wang²

¹*School of Computer Science and Engineering*, ²*School of Civil and Environmental Engineering*

Nanyang Technological University, Singapore

{xiaoyun001, caochu, limo, wangzhiwei}@ntu.edu.sg

Abstract—Service disruptions of rail transit systems become more frequent in the past decades in urban cities like Singapore, due to various reasons such as power failures, signal errors, *etc.* We study and predict the impact of disruptions to transit systems and commuters. This benefits service providers in making both short and long term plans to improve their services. Specifically, we define two metrics, *stay ratio* and *travel delay*, to quantify the impact. To tackle the main challenge of abnormal data scarcity, *i.e.*, only 6 observed disruptions in our one-year data records, we propose to format the problem into a training problem on a feature space relevant to alternative route choices of the commuters. We demonstrate the new feature space corresponds to more similar data distribution among different disruptions, which is beneficial for training more generalisable predictors for future disruptions. We implement and evaluate our approach with a real-world transit card dataset. The result clearly shows that our method outperforms a range of baseline methods.

Index Terms—Service disruption, impact prediction, data scarcity

I. INTRODUCTION

The rapid rail system is the backbone of the public transit systems (PTS) in urban cities. Malfunction of the rail system even in a small region may have ripple effects and significantly impair the PTS. According to our study on Singapore Mass Rapid Transit (MRT) rail system, major disruptions take place due to many reasons including technical faults, extreme weathers, human injuries, *etc.* The journey of thousands or even tens of thousands of commuters may be impaired. Many of them have to quit the PTS and resort to other transportation alternatives (*e.g.*, taxis).

This paper aims at predicting the impact of rail system disruptions at the time of occurrence. Such knowledge not only benefits the PTS provider in understanding the degradation of service, making better emergent plans and planning appropriate new services in PTS to improve system resilience, but also benefits commuters in preparing for the hazards brought by disruptions [1] [2]. Specifically, we define the following two metrics to assess the impact of disruptions. (1) **Stay ratio** indicates the percentage of rail riders who choose to stay within the PTS and take alternative rail lines and/or buses to complete their trip. (2) **Travel delay** indicates the extra time spent on alternative routes for those who stay within the PTS. Obviously, higher stay ratio and lower travel delay indicate smaller impact by a disruption.

Although there have been efforts made to analysing the influence of abnormal conditions of railway on commuters

[3]–[5], most of them apply empirical knowledge or simplified human behaviour models to reason human choices, and based on that analyze the impact on commuters. Some exploit real transportation data to understand human behaviours, but they are often limited to normal PTS conditions. In this paper, taking a unique approach, we explore the transportation data during rail system disruptions and learn from the true human choices. We train a human behaviour model from those abnormal data and apply the model to predict the impact of future disruptions.

Being simple in rationale, our approach is especially challenged due to the scarcity of abnormal data, *i.e.*, those from only 6–8 major disruptions per year. A direct challenge comes from the lack of training data for us to build an accurate model using supervised learning. The limited observation of disruptions makes the trained model difficult to generalize, *i.e.*, applicable to future disruptions unseen in the training stage. The problem becomes more challenging if we consider that only the trips of regular commuters (which is a small portion of the total affected commuters) can be utilized to analyze human behaviours, extract features and label impact metrics, because for irregular commuters there is no way to infer their original travel intention and thus no confidence with regard to their choices under disruptions.

In order to address the above challenges, we propose a novel idea of domain projection to tackle the data distribution mismatch between training and testing sets especially in the situation of data scarcity. Similar but different to the situation of canonical transfer learning, our data in both the training and testing sets is scarce and hence no big picture of the distribution can be profiled. Therefore, we claim the importance of proactively finding a feature space where the training and testing disruptions share similar distributions of extracted features. Specifically, the proposed domain projection method converts the original training problem on the feature space relevant to disruption itself to a new training problem on a different feature space relevant to alternative route choices of the commuters, which unifies our view of disruptions by their effect on commuter route choices. A model trained from the converted feature space can thus be generalized to arbitrary disruptions as long as the commuter route choices can be inferred from the disruptions.

Our contributions are summarized as follows:

- To the best of our knowledge, this is the first study of

impact prediction of rail system disruptions that learns models from true human behaviours in disruptions.

- We propose a novel domain projection method to address the challenges arising from data scarcity, with which we are able to build an accurate and more generalizable model for arbitrary disruptions.
- We implement and experimentally evaluate our approach with the Singapore MRT ride records in year 2015 that involve 6 major disruptions. The results demonstrate that our method outperforms all the baseline methods.

The rest of the paper is organized as follows. Section II presents the definition of the problem. Section III to IV detail our method for impact prediction. We present evaluation settings and results in Section V. We review the related work in Section VI and conclude this paper in Section VII.

II. PROBLEM DEFINITION

In this section, we define the problem of impact prediction of disruptions, and present an overview of our methodology.

We treat a rail network as a directed graph $G = (V, E)$, where V represents the set of stations and E the directed rail links between stations. A disruption results in some link(s) removed from the graph. Each disruption lasts for a period of time which is usually unknown at the time of occurrence. We only consider *major disruptions* which last longer than 30 minutes. A formal definition of disruption is given as follows.

Definition 1 (Disruption). A disruption refers to a period of no train service on a set of adjacent links of G . A disruption $e^{(i)}$, $i = 1, 2, \dots$, is represented by a tuple $(T^{(i)}, G^{(i)})$, where $T^{(i)}$ is the starting time and $G^{(i)} = (V, E^{(i)})$ is the disrupted rail network ($E^{(i)} \subseteq E$) with $E \setminus E^{(i)}$ removed links.

Due to the disruption, stranded commuters either stay in the PTS (*e.g.*, wait for system resumption, or choose alternative routes from bus network and the remaining rail network) or leave the PTS and look for other transportation modes (*e.g.*, taxis). We use Voronoi Diagram [6] to partition the city into Voronoi cells centering at rail stations, where each cell is a *region* containing the rail station and nearby bus stations. We can then describe each of the commuter trip as an OD (Origin-Destination) sample between any two of those regions. Then, we define *affected* ODs formally as below.

Definition 2 (Affected OD). During a disruption $e^{(i)}$, an affected OD is a pair of stations (u, v) that is unreachable in $G^{(i)}$, or is tortuous, *i.e.*, $d^{(i)}(u, v) - d(u, v) > \lambda$, where $d(u, v)$ and $d^{(i)}(u, v)$ are the topological distances between u and v respectively in G and $G^{(i)}$. λ is a scalar set to be sufficiently large (*e.g.*, $\lambda = 10$).

Affected commuters of a disruption are those with their OD being one of the affected ODs of that disruption. In this paper, we aim at training two impact predictors separately for two impact metrics, namely, stay ratio and travel delay. Formally, given an affected OD (u, v) of a disruption $e^{(i)}$, the stay ratio $I_s^{(i)}(u, v)$ and travel delay $I_t^{(i)}(u, v)$ are defined as

$$I_s^{(i)}(u, v) = \tilde{r}_{uv}^{(i)} / r_{uv}^{(i)}, \quad (1)$$

$$I_t^{(i)}(u, v) = \tilde{t}_{uv}^{(i)} - t_{uv}^{(i)}. \quad (2)$$

where $r_{uv}^{(i)}$, $t_{uv}^{(i)}$ are the normal ridership and travel time on the original (rail) route, and $\tilde{r}_{uv}^{(i)}$, $\tilde{t}_{uv}^{(i)}$ are the disruption-affected ridership and (averaged) travel time on alternative routes in PTS, between (u, v) during the disrupted period of the day.

We were granted access to over one-year transit card records of bus and rail rides in Singapore from June 2015 to June 2016. Each record contains user ID, boarding and alighting stations and timestamps, as well as the bus/rail service name. We also obtained the information of historical disruptions from Singapore MRT operators from their official Twitter announcements that contains disruption date, starting time, ending time, and location (*i.e.*, rail line and the stretch of disrupted stations). We specifically study those data on working days.

To train the impact predictors, we first propose the domain projection to convert the prediction problem in the domain of disruption into that in the domain of interested alternative routes (IARs) that may be chosen by the commuters during disruptions, where we may address the challenge of data scarcity and train a generalizable model (Section III). Regular commuters are identified from historical trips, whose travel patterns under normal conditions are stable and their choices under disruptions will be utilized to label impact metrics and to construct impact predictors. After that, we generate IARs and construct impact predictors using light-weight machine learning techniques (Section IV). For a given affected OD in a given disruption, the two predictors can give anticipated stay ratio and travel delay as a result of the impact.

III. DOMAIN PROJECTION

In this section, we describe the idea of domain projection, which attempts to find a feature domain where the trained model is more generalizable to future disruptions.

Impact metrics are related to features of disruption and affected OD. For example, a disruption with more broken links may lead to a lower stay ratio due to the greater mismatch in a sudden between transit demand and capacity, and a larger travel delay due to longer travelling time on alternative routes in PTS. A straightforward solution is to train a model via supervised learning based on disruption and OD features such as starting time, number of broken links, *etc.* Such a method, however, may result in under-fitting since the model being trained on scarce training disruptions may not capture the functional relation, between features and impact metrics, that can extend to future disruptions. To illustrate such a view, we name the domain of disruption and OD features as $\mathbb{D}_1 = (\mathcal{X}_1, P_1)$, where \mathcal{X}_1 is a d_1 -dimension feature space, and each affected OD can be represented by a point $X = \{x_1, \dots, x_{d_1}\} \in \mathcal{X}_1$ with the probability denoted by $P_1(X)$. We visualize affected ODs of observed disruptions as points in \mathcal{X}_1 (features listed in Table II), using PCA (principal component analysis) to reduce them to planar points. From the result plotted in Figure 1(a), we see that distributions of points of different disruptions hardly coincide. Essentially, such a problem results from the limited number of disruptions that we can observe - the data distribution of what we can observe and the data distribution of what we want to predict mismatch.

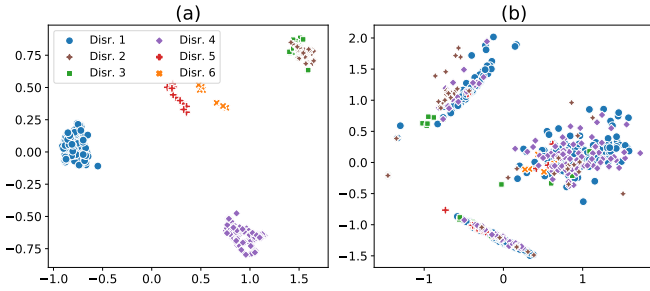


Fig. 1: Visualizing points of affected ODs of 6 observed disruptions in (a) domain \mathbb{D}_1 and (b) domain \mathbb{D}_2 .

We propose to translate the prediction problem in domain \mathbb{D}_1 to a new domain \mathbb{D}_2 , where the feature space has the following property: all observed disruptions should contain points of similar distribution. This guarantees the model applicability to future disruptions potentially with similar distribution. In our case, we aim at describing an disruption around the features of IARs (*interested alternative routes*), which are the alternative routes in PTS likely to be chosen by affected commuters. For an affected OD, the features of available IARs influence the choices of affected commuters, because they concern whether the IARs satisfy their requirements, *e.g.*, a short waiting time, no transfers, *etc.* The travel time of an IAR is also closely correlated to the travel delay. Therefore, we denote the new domain as $\mathbb{D}_2 = (\mathcal{X}_2, P_2)$, where \mathcal{X}_2 is a d_2 -dimension feature space of IAR features (listed in Table III that are categorized into service-wise and route-wise features). Each affected OD can be represented by a point $X = \{x_1, \dots, x_{d_2}\} \in \mathcal{X}_2$ with the probability denoted by $P_2(X)$. For comparison, we visualize affected ODs of observed disruptions as points in \mathcal{X}_2 in the same way as what we do in \mathcal{X}_1 . We can see from Figure 1(b), that the clusters of points are mixed across disruptions, suggesting highly overlapped distributions in the feature space \mathcal{X}_2 .

IV. IMPACT PREDICTION

In this section, we first show how to identify regular commuters and analyse their choices under disruptions. We then describe how we generate interested alternative routes (IARs), based on which we build impact predictors.

A. Obtaining Regular Commuters

Affected commuters have three choices: taking alternative routes in PTS, taking other transportation modes, or mixing of both. From the transit card records, we are able to find their traces in PTS during disruption. However, as travel plans of the commuters might change day by day, rendering the difficulty of deciding their original OD, it is non-trivial to infer affected commuters' choices. We propose to identify regular commuters whose travel behaviors (*i.e.*, departure time and OD) are relatively stable and thus their original OD can be determined. After that, we focus our study on those regular commuters on behalf of all PTS users.

To identify regular commuters for a specific disruption, we obtain the list of ODs for each commuter during the disruptive hours of the day during a sufficiently long past period (*e.g.*, pass two months) before the date of disruption. Then the frequencies of distinct ODs in the list are denoted by m_1, m_2, \dots , and the highest frequency is denoted by m_h . We define the OD with frequency m_h as the *dominant OD*. Then we start to filter out commuters that are irregular. We remove those whose $m_h < \max(\frac{1}{2} \sum_i m_i, \epsilon)$, to filter out occasional riders to the rail system and ϵ is set to 5. We then cluster commuters with the same m_h value together. In each cluster, we remove top-25% commuters with the largest *entropy*, *i.e.*, $-\sum_i q_i \log q_i$, where $q_i = m_i / \sum_i m_i$. We assume dominant OD as the regular commuter's original travel plan, and consider only affected regular commuters whose dominant OD is one of the affected ODs.

B. Analysing Under-disruption Choices

Given the set of affected regular commuters, we further analyse their choices under disruptions. We find that over 90% of affected ODs have a low stay ratio which is less than 0.5, indicating that affected commuters are prone to leave the PTS when a disruption occurs. For those who stay in PTS until arrival, over 90% have less than 50 minutes travel delay.

In addition, we analyse the alternative routes chosen by commuters who stay in PTS until arrival. A route is represented by a list of bus and/or rail services, each of which is indicated by the bus service name or a pair of rail stations. The corresponding list of transit modals for an alternative route is regarded as its *route pattern*. We find that the top-4 route patterns chosen by affected commuters, namely (bus, bus), (bus, rail) and (rail, bus), account for a percentage of 86%. Besides, the walking distance between two successive services is usually within 500 meters, and the detouring rate (*i.e.*, the ratio of the distance of alternative route to that of the original route) is less than 1.5. These will serve as constraints when we generate IARs in Section IV-C.

C. Interested Alternative Routes Generation

For each affected OD, we first generate candidates by routing on the network of bus and disrupted rail systems using depth-first searching [7]. Candidates should satisfy constraints about walking distance (< 500 meters), detouring rate (< 1.5) and route pattern (being in top 7 patterns) mentioned in Section IV-B. We also gather real IARs from our transit records. We label real IARs by 1, and negative candidates that are not real IARs by 0. The number of real IARs is extremely limited (*i.e.*, around 2.6 per affected OD) while the number of candidates is huge, leading to a significant imbalance (about 1 to 15,000) between the two classes. To alleviate the imbalance, we conduct negative sampling on the set of candidates. That is, to sample instances of candidates according to their similarity to real IARs, *i.e.*, the more similar an instance is to any of the real IARs, the higher probability it is being sampled. To be specific, we consider several dimensions for similarity, namely, the numbers of service transfers, rail stations and bus stations,

as well as the length of waiting time and the distance of walking. Each dimension is normalized into the range of $[0, 1]$. Then each candidate is represented by a 5D vector, and the similarity between any pair, one from real IARs and the other from candidates, is calculated by cosine similarity.

We use the features in Table III as well as the labels from real IARs and sampled negative candidates to train a binary classifier for IAR identification, so that for any future disruption we can use it to distinguish real IARs from others. We train the classifier using *ensemble learning*, a supervised learning method that trains a couple of models (*i.e.*, Decision Tree in our case) using different subsets of training data, and aggregates the results by majority voting. It also restrains the class imbalance problem. For any affected OD of future disruptions, we can first generate candidate IARs and then identify real IARs using the classifier.

D. Predictors Building

We train the stay ratio and travel delay predictors separately. Given an affected OD, we leverage Equation (1) and (2) to calculate stay ratio and travel delay (*i.e.*, labels). And for the input, the features of each IAR are concatenated to form a vector (inapplicable features are filled by zeros). Each IAR belongs to one group according to the length of route pattern. In each group, element-wise statistical aggregations, namely, *mean*, *max* and *min*, over all group members are calculated and we append the result together to form the aggregated vector. We then concatenate the aggregated vectors of all groups (in ascending order by the length of route pattern), to form a new feature vector, where *backward elimination* [8] is conducted to select final vital features. With the training samples of processed IAR features and labels we apply SVR to model the relationship between the IAR features and two impact metrics, and train the two predictors.

V. EVALUATION

A. Experimental Setup

We obtain the information of disruptions (*e.g.*, time and locations) from instant tweets posted by MRT operators. We finally get 6 major disruptions between June 2015 to June 2016, detailed information of which are summarized in Table I. In addition, to identify regular commuters and generate alternative routes, PTS static metadata (*i.e.*, bus/rail routes, station ids and locations) are obtained from the LTA Online Datamall [9], and walking/travel distance between geographical locations are acquired via Google’s Direction APIs.

We conduct experiments for the following methods:

- \mathbb{D}_1 -SVR: which is built on original problem domain \mathbb{D}_1 with SVR applied to features of disruption and affected OD listed in Table II. Backward elimination is used before training the model.
- PIRD: **P**redicting **I**mpact of **R**ail **D**isruptions, the method proposed in this paper, using IAR features in Table III.
- \mathbb{D}_2 -Oracle: which is built on domain \mathbb{D}_2 and is implemented the same way as PIRD, but with feature input from real IARs (which are not available during

TABLE I: Summary of the 6 disruptions.

	date	start time	duration (min)	affected line(s)	#links re-moved	#affected commuters	#regular commuters
1	20150707	19:30	110	EW,NS	108	330,000	130,000
2	20151013	08:00	30	NE	12	38,000	33,000
3	20151026	05:25	90	NE	30	52,000	24,000
4	20151125	05:50	140	NS	8	79,000	33,000
5	20151217	19:50	135	EW	4	17,000	6,000
6	20160322	11:10	160	EW	6	38,000	4,000

prediction and can only be obtained after the commuters complete their trips). This approach utilizes practically not available information and its performance represents the upper-bound for comparison.

- PIRD-LR: which is implemented the same way as PIRD, except that the final regression model applied to IAR features is linear regression instead of SVR.
- \mathbb{D}_2 -SVDD: which implements the same approach as PIRD except that it selects candidate IARs using SVDD [10] instead of ensemble learning. Using OCC (only one class) of positive samples avoids the imbalance issue but loses information from negative samples.

We adopt a *leave-one-out* scheme to evaluate the proposed impact predictors. Each time, we take OD samples from 5 disruptions as training set and OD samples in the remained disruption as testing set, representing the same setting when we apply our solution in reality, *i.e.*, we have historical disruptions to train predictors for a future disruption. The hyper-parameters are tuned using 5-fold cross validation on the training set. We run the experiment for each of the evaluated methods 200 rounds. We use MAE (mean absolute error) to evaluate the performance averaged over all 6 events. For each method, we calculate an *average* MAE and a *worst* MAE over all tested ODs. Specifically, we denote the absolute error of the j -th OD in the k -th round when using the i -th disruption as test data as $AE_{jk}^{(i)}$, for $k = 1, \dots, 200$, $i = 1, \dots, 6$ and $j = 1, \dots, J_i$, where J_i is the number of affected ODs tested in the i -th disruption. Then the *average* MAE is derived as $\sum_i \sum_k \sum_{j=1}^{J_i} AE_{jk}^{(i)} / (200 \sum_i J_i)$ and the *worst* MAE is derived as $Max_i (\sum_k \sum_{j=1}^{J_i} AE_{jk}^{(i)} / (200 J_i))$.

B. Prediction Accuracy

Both the *average* and *worst* MAE of evaluated methods on stay ratio and travel delay prediction are provided in Table IV. For stay ratio prediction, PIRD provides 0.11 *average* and 0.12 *worst* MAE, and for travel delay prediction, PIRD gives 11.9 minutes *average* and 14.5 minutes *worst* MAE. PIRD outperforms \mathbb{D}_1 -SVR. The results suggest the training performance over domain \mathbb{D}_2 outperforms that conducted on \mathbb{D}_1 . PIRD achieves MAEs close to that of \mathbb{D}_2 -Oracle in stay ratio prediction, which indicates our method works well in producing real choices of commuters. The results also suggest PIRD has close *average* MAEs to PIRD-LR’s in both stay ratio and travel delay prediction, but outperforms PIRD-LR in the *worst* MAE, which suggests the performance gain of applying SVR over linear regression.

TABLE II: Features of disruption and affected OD.

No.	Feature Description
1	starting time, <i>i.e.</i> , T (each time slot of 0.5h)
2	# of removed links
3	# of affected ODs
4,5,6,7	binary value for a rail line to indicate normal/disruptive state
8,9	location (<i>i.e.</i> , latitude and longitude) of origin station
10,11	location of destination station
12,13	the rail line where the origin and dest. station located
14	# of affected stations on the rail route between affected OD
15	# of stations from origin station to the nearest affected station
16	# of stations from dest. station to the nearest affected station

TABLE IV: MAE comparison of different methods.

Impact Metric	\mathbb{D}_1 -SVR	PIRD-LR	PIRD	\mathbb{D}_2 -Oracle
<i>avg.</i> stay ratio/ travel delay(min)	0.16/13.8	0.11 /11.4	0.11 /11.9	0.11 /10.2
<i>worst</i> stay ratio/ travel delay(min)	0.22/15.5	0.14/15.1	0.12 /14.5	0.14/ 11.2

TABLE V: Performance comparison with variant approach.

	stay ratio <i>avg./worst</i> MAE	travel delay <i>avg./worst</i> MAE	% of candidates
\mathbb{D}_2 -SVDD	0.13/0.14	11.8 /14.9 min	8.6
PIRD	0.11 /0.12	11.9/ 14.5 min	4.4

Figure 2(a) and (b) present statistics of prediction errors for different ODs across different disruptions. We have two observations from these two figures. First, PIRD performs close to \mathbb{D}_2 -Oracle for most disruptions and may even outperforms it (*e.g.*, for the Disruption 6 due to the fact that there are insufficient regular commuters and very few real IARs that can be used by \mathbb{D}_2 -Oracle). Second, PIRD performs well in generalization. In the data records, Disruption 1 has 354 affected ODs, which account for 60% of the total number of ODs in the study. When we use Disruption 2-6 as training data, and build the model to predict the impact to ODs in Disruption 1, however, we see that the error (*i.e.*, an average of 0.1 for stay ratio and 9 minutes for travel delay) is not apparently higher than what we can obtain for other disruptions. It suggests that PIRD is able to capture critical features and has strong ability in generalization with small training data.

We also compare the performance of PIRD and that of \mathbb{D}_2 -SVDD. Table V summarizes the *average* and *worst* MAEs derived over 6 disruptions (based on the same procedure as adopted for Table IV). Specifically, for stay ratio, PIRD provides smaller *average* and *worst* MAEs (0.11 and 0.12) than \mathbb{D}_2 -SVDD (0.13 and 0.14). For travel delay, PIRD achieves very close performance to \mathbb{D}_2 -SVDD but our method leverages much less data for training, accounting for only 4.4% of all IAR candidates while \mathbb{D}_2 -SVDD’s training data usage nearly doubles our data. The results indicate that the ensemble learning design in PIRD can effectively select meaningful IARs that are possibly chosen by commuters during a disruption.

C. Stability

Stability provides tolerance to perturbations from training data, which is very important to our problem as perturbations

TABLE III: Features of IAR.

Type	No.	Feature Description
service -wise	1,2	the shortest and longest waiting time
	3	walking distance
	4	binary value for transit modal (bus/rail)
	5	# of bus/rail stations traveled
	6,7	normal boarding/alighting ridership of the service around T
route -wise	8	# of other bus services with nearby (<500m) bus stations
	1	# of bus/rail services
	2	normal ridership of the first service
	3	# of rail stations travelled in the original route

may come from data noises, variation from new disruptions, emergent actions taken during the disruptions, and so on. We evaluate the stability of PIRD in comparison with other methods. Stability can be reflected from the results of predictors. Figure 2(c) and (d) present the distributions of MAE (averaged across 6 disruption) from each round of prediction after zero-mean normalization, where each line represents 200 MAE points. The sharper the ascending curve is, the more consistent the MAEs are, indicating a stabler output. Figure 2(c) shows that PIRD is stabler than \mathbb{D}_1 -SVR and \mathbb{D}_2 -SVDD, and is close to the performance of \mathbb{D}_2 -Oracle. With regard to the travel delay in Figure 2(d), PIRD is stabler than other methods, and even outperforms \mathbb{D}_2 -Oracle mainly due to the fact that \mathbb{D}_2 -Oracle has insufficient real IARs for training stable models. The results suggest the high stability of PIRD as compared with other alternatives.

VI. RELATED WORK

Impact Prediction. There have been efforts made to predict the impact of transportation incidents, *e.g.*, railway disruption, traffic accidents, *etc.* Some works predict the impact by reasoning human reactions or the damage to network structure, most of which lack measurement study of real incidents. For examples, Sun *et al.* [5] estimates the normal spatio-temporal distribution of commuters in rail system, and try to infer the number of affected commuters when there is a disruption. Sun *et al.* [4] try to reason commuters’ travel delay based on their choices (*e.g.*, stay or leave PTS). Yin *et al.* [11] define the impact as the damage to rail network efficiency, and utilize graph theory to quantify the impact of disruption. Some works predict impact based on actual mobility data measured from real world. Examples include Pan *et al.* [12] who take the average impact of similar historical incidents to predict that of future incidents, Fang *et al.* [3] who leverage contextual features and post-incident travel delays to predict future travel delays, and Garib *et al.* [13] who use statistical models based on contextual features to predict travel delay. Most existing studies have not thoroughly investigated the ability of generalization and are not validated with real world incidents at the scale of this paper. Other studies focus on forecasting the traffic flow under anomalous conditions [14]–[16] taking a period of post-incident traffic flow as input. The traffic flows, however, cannot be translated to fine-grained impact to commuters. To sum up, so far there is no existing study




Fig. 2: (a) and (b): statistics of prediction errors for different disruptions' ODs; (c) and (d) stability comparison between evaluated methods for stay ratio (a) and travel delay (b).

which measures impact from real incidents, and meanwhile explores the model generalizing ability to predict the impact of a variety of future incidents.

Feature manipulation. The proposed domain projection method shares similarity with some other feature manipulation techniques, *e.g.*, *feature engineering* [17]. The data distribution mismatch between training and testing sets however is not the focus of feature engineering. The most relevant study is that of *transductive transfer learning* [18], which transfers knowledge from the training set to the testing set when the data distribution of their feature spaces are different. Existing studies such as Zadrozny *et al.* [19] and Daume III *et al.* [20] use re-sampling or statistical adaptation of the distribution of training set to that of the testing set. All existing studies assume ample labeled testing data, and the global distribution of population can be profiled, which is distinct from our case.

VII. CONCLUSION

We propose a comprehensive solution to predict the impact of rail system disruptions, based on the real behaviors of affected commuters during disruptions. To tackle the challenge of training data scarcity, We propose to project a disruption and its affected OD into a different domain of features abstracted from commuters' alternative route choices. The training accuracy and generalizing ability are greatly improved. Experimental results using real-world data demonstrate the effectiveness of our proposed solution.

ACKNOWLEDGMENTS

This research is partially supported by NRF Singapore under its grant SDSC-2019-001, Alibaba Group through Alibaba Innovative Research (AIR) Program and Alibaba-NTU Singapore Joint Research Institute (JRI), and Singapore MOE Tier 1 grant RG18/20. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of funding agencies. We thank Singapore Land Transport Authority (LTA) for kindly sharing with us the transit card data to study with.

REFERENCES

- [1] P. Zhou, Y. Zheng, and M. Li, "How long to wait? predicting bus arrival time with mobile phone based participatory sensing," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*, 2012, pp. 379–392.
- [2] Z. Liu, Z. Gong, J. Li, and K. Wu, "Mobility-aware dynamic taxi ridesharing," in *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 2020, pp. 961–972.
- [3] Z. Fang, Y. Yang, D. Zhang *et al.*, "Mac: Measuring the impacts of anomalies on travel time of multiple transportation systems," *Proceedings of ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 2, p. 42, 2019.
- [4] H. Sun, J. Wu, L. Wu, X. Yan, and Z. Gao, "Estimating the influence of common disruptions on urban rail transit networks," *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 62–75, 2016.
- [5] L. Sun, D.-H. Lee, A. Erath, and X. Huang, "Using smart card data to extract passenger's spatio-temporal density and train's trajectory of mrt system," in *Proceedings of the ACM SIGKDD international workshop on urban computing*. ACM, 2012, pp. 142–148.
- [6] G. Voronoi, "Nouvelles applications des paramètres continus à la théorie des formes quadratiques. deuxième mémoire. recherches sur les parallélogrammes primitifs." *Journal für die reine und angewandte Mathematik*, vol. 134, pp. 198–287, 1908.
- [7] S. Robert, "Algorithms in c, part 5: Graph algorithms," 2002.
- [8] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *Journal of machine learning research*, vol. 3, no. Mar, pp. 1157–1182, 2003.
- [9] S. L. T. A. Datamall. <https://www.mytransport.sg/content/mytransport/home.html>.
- [10] D. M. Tax and R. P. Duin, "Support vector data description," *Machine learning*, vol. 54, no. 1, pp. 45–66, 2004.
- [11] H. Yin, B. Han, D. Li *et al.*, "Evaluating disruption in rail transit network: a case study of beijing subway," *Procedia Engineering*, vol. 137, pp. 49–58, 2016.
- [12] B. Pan, U. Demiryurek, C. Shahabi, and C. Gupta, "Forecasting spatiotemporal impact of traffic incidents on road networks," in *2013 IEEE 13th International Conf. on Data Mining*. IEEE, 2013, pp. 587–596.
- [13] A. Garib, A. Radwan, and H. Al-Deek, "Estimating magnitude and duration of incident delays," *Journal of Transportation Engineering*, vol. 123, no. 6, pp. 459–466, 1997.
- [14] R. Yu, Y. Li, C. Shahabi, U. Demiryurek, and Y. Liu, "Deep learning: A generic approach for extreme condition traffic forecasting," in *Proceedings of SIAM*, 2017, pp. 777–785.
- [15] D. Ding, M. Zhang, X. Pan, M. Yang, and X. He, "Modeling extreme events in time series prediction," in *Proceedings of the 25th ACM SIGKDD*, 2019, pp. 1114–1122.
- [16] M. Castro-Neto, Y.-S. Jeong, M.-K. Jeong, and L. D. Han, "Online-svr for short-term traffic flow prediction under typical and atypical traffic conditions," *Expert systems with applications*, vol. 36, no. 3, pp. 6164–6173, 2009.
- [17] S. Khalid, T. Khalil, and S. Nasreen, "A survey of feature selection and feature extraction techniques in machine learning," in *2014 Science and Information Conference*. IEEE, 2014, pp. 372–378.
- [18] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE TKDE*, vol. 22, no. 10, pp. 1345–1359, 2009.
- [19] B. Zadrozny, "Learning and evaluating classifiers under sample selection bias," in *Proceedings of ICML*, 2004, p. 114.
- [20] H. Daume III and D. Marcu, "Domain adaptation for statistical classifiers," *Journal of artificial Intelligence research*, vol. 26, pp. 101–126, 2006.