PARAMETER SELECTION FOR "PCMA*" USING SURROGATE-ASSISTED BLACK-BOX OPTIMIZATION AND MICROSCOPIC TRAFFIC SIMULATION

Bernhard Werth[(a)], Erik Pitzer[(b)], Michael Affenzeller[(c)], Christian Backfrieder [(d)]

[(a),(b),(c)] Heuristic and Evolutionary Algorithms Laboratory, University of Applied Sciences Upper Austria, Hagenberg
[(a),(c)] Institute for Formal Models and Verification, Johannes Kepler University, Linz, Austria
[(d)] Research Group Networks and Mobility, University of Applied Sciences Upper Austria, Hagenberg

[(a)]bernhard.werth@fh-hagenberg.at, [(b)]erik.pitzer@fh-hagenberg.at, [(c)]michael.affenzeller@fh-hagenberg.at
[(d)]christian.backfrieder@fh-hagenberg.at

ABSTRACT
Routing algorithms have become significantly more sophisticated in recent years as increasingly more data on road networks, vehicles and drivers becomes available. The correct parametrization of such algorithms is a difficult task since the underlying traffic model needs to reflect the complexity of the algorithm and should resemble reality closely. A major drawback of such models is that evaluating a certain parameter setting may become computationally expensive, preventing the use of conventional optimization algorithms. In this work a combination of surrogate assisted black box optimization and microscopic traffic simulation is used to optimize the parameters of a recently published routing algorithm.

Keywords: traffic simulation, surrogate assisted optimization, evolutionary algorithms, noisy optimization

1. INTRODUCTION

With the ever increasing number of vehicles especially in urban areas, the issue of traffic congestion becomes more severe. The Urban Mobility Report (Schrank, Eisele, and Lomax 2012) emphasizes the need for new solutions in battling the growing costs created by traffic jams in US American cities. Microscopic traffic simulation could help finding such solutions as it is a powerful tool to evaluate the impacts of many different influences on a road network. Traffic light controls, the percentage of automatic vehicles, probabilities of accidents or different routing algorithms could be examples of such influences.
Recently a new predictive routing algorithm called *"Predictive Congestion Minimization in Combination with an A*-based router"* (PCMA*) which is based on predictions about bottlenecks and traffic jams in road networks has been proposed (Backfrieder et al 2017). In their experiments the authors noted that correct calibration of the thresholds used in their algorithm is a rather complex task and might even be impossible in a global manner as these thresholds and parameters depend on the specific road network considered as well as on the distribution of vehicles and their start and end points. Since mathematical representations for microscopic traffic simulations are usually not available, the calibration of such parameters on a specific scenario has to be treated as a black-box optimization problem. However, the use of conventional optimization algorithms like *Genetic Algorithms, Particle Swarm Optimization* or *Tabu-Search* is hindered by the fact that they require a large number of function evaluations to converge, which is infeasible when a single simulation run for testing one parameter configuration takes several minutes or even hours depending on the scenario and the degree of parallelization. The use of cheap surrogate models to facilitate the search is therefore imperative.
The remainder of this work is structured as follows: Section 2 provides an overview of surrogate-assisted optimization techniques and which of those are of relevance in the scenario at hand. Section 3 describes the different classes of routing algorithms and very briefly outlines the characteristics of the PCMA* algorithm. In Section 4 the optimization problem with its parameters and bounds as well as the algorithm and surrogate model type are described. Section 5 contains computational results, followed by conclusions in Section 6.

2. SURROGATE-ASSISTED OPTIMIZATION

Using surrogate models in optimization tasks is an established method and often used when the evaluation of solution candidates is either time consuming or associated with additional costs. Surrogate-assisted optimization has been successfully applied to various simulation scenarios including Navier-Stokes flow solvers for aerodynamic simulations (Han 2013), finite element simulations for fluid dynamics (Forrester and Keane 2007), flowsheet simulations in manufacturing environments (Boukouvala and Ierapetritou 2013) and mesoscopic traffic simulation (He 2014).
The main idea of surrogate-assisted optimization is to use interpolation and statistical models that are considerably cheaper to evaluate than performing a full simulation. Therefore, surrogate-based optimization strategies should be able to find good solutions within a considerably smaller number of simulations than conventional optimization algorithms would require. Much alike the wide variety of application scenarios for surrogate-assisted optimization, a considerable number

of algorithms, algorithm extensions and modifications exist. *Particle Swarm Optimization* (Sun et al. 2015), various versions of *Evolutionary Strategies* (Loshchilov, Schoenauer, and Sebag2013) and even interactive *Genetic Algorithms* have been enhanced with surrogate models.

In this work we decided to focus on an adaptive sampling scheme similar to the one described in (Wang and Shan 2007), depicted in Figure 1. First, an initial set of samples is created and evaluated. In the next step a Gaussian process model is constructed, which can then be used to select a solution candidate that is optimal with respect to some infill criterion. This candidate is then evaluated on the expensive simulation. After obtaining a result value from the simulation, the process can be started anew until a certain computational budget is exhausted.

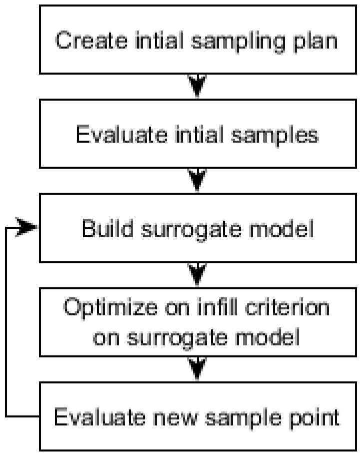

Figure 1: Adaptive Sampling Scheme

A variety of different types of models have been used in this context, with *polynomial regression*, *neural networks*, *support vector machines*, *radial basis function regression* and *Gaussian process regression* being among the most popular. Gaussian processes have the additional benefit of providing a measure for the models own uncertainty, which has been used in the *expected improvement* infill criterion in the widely used "Efficient Global Optimization" algorithm (EGO) by (Jones, Schonlau, and Welch 1998). The original expected improvement criterion was designed for deterministic black box functions without noise. Several new infill criteria for noisy optimization have since been proposed. A comprehensive overview and a benchmark comparison of such criteria for noisy optimization are given in (Picheny, Wagner, and Ginsbourger 2013). They are of special interest for this work since traffic systems are inherently influenced by many random effects ranging from drivers characteristics like attention span, reaction time, or willingness to take risks to outer influences like weather, slight variations in departure times or technical failures of vehicles and traffic control systems.

3. A PRIMER ON ROUTING ALGORITHMS

Nowadays, camera systems, traffic flow sensors, traffic control authorities and public transport agencies continuously produce large quantities of data concerning the condition of roads, vehicles, traffic lights and other infrastructure. Routing algorithms and automatic traffic guidance systems have evolved with this increased development in traffic surveillance infrastructure and can be roughly divided into three categories:

Dynamic Traffic Assignment systems incorporate theoretical traffic optimization solutions with the goal of achieving a suitable user equilibrium where no vehicle can achieve a shorter time to reach its destination by changing its route (Janson 1991). Major drawbacks of these methods are the lack of ability to react to unforeseen congestions and a high computational effort.

Reactive Routing Algorithms take the current traffic situation and condition of the road network into account to dynamically change the routes of vehicles in the vicinity of current congestion areas (Kesting, Treiber, and Helbing 2010). Unfortunately, some reactive algorithms provide the very same alternative route to all vehicles in a system, thereby effectively moving the congestion from one part of the road network to another rather than actually solving the problem.

Predictive Routing Algorithms use historical data to predict future congestions via a variety of models, which limited their application to highways or higher level interurban roads, as data for lower level roads was scarcely available. Recent developments like the rising number of trackable smart devices, may change this situation in favor of predictive routing systems.

The comparison of such algorithms can be done on traffic simulations with different levels of abstraction. In this work we chose microscopic traffic simulation as our means to evaluate and compare different algorithm parametrizations, as microscopic simulation not only captures reality more closely than a macroscopic model, but also permits to model aspects of reality that are not considered by the algorithm.

The PCMA* algorithm (Backfrieder et al 2017) uses both a predictive and a reactive routing strategy. It combines a routing algorithm that takes the current situation of the road network into account by weighting edges in the routing graph via the *Speed Average* or *Greenshield's* method (Backfrieder et al 2017, Greenshields 1935) with a congestion prediction method that is based around the calculation of so called "footprint" values for individual intersection. These footprints are influenced by information stemming from vehicular communication.

As rerouting all vehicles simultaneously just shifts congestions from one place to another, no vehicle should be treated exceptionally unfavorably and permanent rerouting of the same vehicle might upset the driver, PCMA* employs a strategy for selecting which vehicles to reroute when a congestion is either predicted or dynamically detected.

The predictive and reactive parts of the algorithm and the selection of rerouted vehicles use thresholds to decide when a congestion occurs or a vehicle should be exempt from further rerouting maneuvers. The predictive component additionally has parameters specifying how far into the future predictions are considered relevant. These parameters need to be fine-tuned to the specific scenarios and road networks in order to achieve optimal results.

4. COMPUTATIONAL EXPERIMENTS

The goal of the optimization task will always be the minimization of the cumulative time spent by all vehicles on the road network. The seven real-valued parameters listed in Table 1 should be minimized and therefore constitute a solution candidate. Evaluation is performed via a full simulation in the microscopic traffic simulation framework TraffSim (Backfrieder, Mecklenbräuker, and Ostermayer 2012).

Table 1: Parameter ranges for the PCMA*

Name	Min	Max
Congestion level	1	15
Speed average congestion threshold	0	1
Rerouting threshold distance (reactive)	0	10
Rerouting threshold distance (predictive)	0	10
Congestion threshold (Vehicles per second)	0.1	3
Footprint prediction time resolution	5000	60000
Footprint prediction maximal forecast	1200000	2600000

Congestion level is relevant for the reactive routing component only and defined as the approximate size of a traffic congestion and all vehicles within this many road segments will be subjected to rerouting.

When the mean speed of vehicles passing a traffic node is below the *speed average congestion threshold* a congestion is detected by the reactive part of the algorithm.

It is unwanted for a vehicle to switch to an unreasonably longer alternative route as this would be perceived as an unfair treatment. If the ratio of the new route's length to the old route's length is larger than the rerouting threshold distance, the alternative route is no longer considered feasible. This threshold can differ for the reactive and the predictive parts of the algorithm.

When the predicted number of vehicles that should pass a certain node of the traffic graph exceeds the threshold specified by the *Congestion threshold* the predictive component will predict a congestion.

Footprint prediction time resolution and *Footprint prediction maximal* forecast specify the time frames for footprint calculation and the limit at which point predictions into the future are no longer considered relevant. Both parameters only concern the predictive part of the PCMA*.

Tests are performed on two different road networks. One is a partial network of the northern part of the city of Linz and the other one is artificially generated with the goal to resemble American downtowns (Lindorfer et al. 2013). Start and end nodes of the vehicles are chosen so that most of the vehicles follow a general direction through the network as it would happen in a morning traffic peak where more traffic happens towards the industrial areas of a city.

As stated in Section 2, the infill criteria created for noisy black box optimization are especially interesting for the task at hand. Table 2 lists the infill criteria used in our experiments.

Table 2: Tested Infill Criteria

Name	Published by
(noiseless) Expected Improvement (EI)	Jones, Schonlau, Welch (1998)
Augmented Expected Improvement (AEI)	Huang et al. (2006)
Plugin Expected Improvement(PEI)	Osborne, Garnett, Roberts (2009)
Expected Quantile Improvement(EQI)	Picheny, Ginsbourger, Richet (2010)
Minimal Quantile Criterion(MQC)	Cox, John (1992)
Expected Quality(EQ)	
Neighbor Distance (ND)	

Additionally to the proposed criteria, the strategy of using the model prediction directly as an infill criterion (Expected Quality) and the approach of maximizing the distance to the nearest neighbor (Neighbor Distance) are used to generate baselines for maximal exploitation (EQ) and maximal exploration (ND) are employed.

The shape of the road network might not be the only factor influencing the optimal parameter selection for the PCMA*, therefore, the number of vehicles in each scenario is varied to be 500, 1000, 1500 and 2000. Also the distribution of departure times might have a considerable impact on the performance of collaborative routing algorithms which is why we chose two different distributions: a uniform distribution with start times in [t, t+30 min] and a normal distribution with $\mu = t$ and $\sigma = 15$ min. Since neither weather nor time of day are considered within the simulation the parameter t has no influence on the results and is set to an arbitrary but fixed value.

Initialization: For creating the initial sampling plan a set of 50 points that are almost-optimal with respect to the well-known maximin criterion (Johnson 1990), which is defined as maximizing the minimal distance between all points of a set, is created using an evolutionary strategy (Beyer and Schwefel 2002) with $\mu = 20$ and $\lambda = 70$ and 500 generations.

Model Building: The surrogate model used is a Gaussian process regression with a constant mean function and a rational quadratic covariance kernel that assumes a homoscedastic error. The hyper parameters of each new model are selected by optimizing the log-likelihood with a gradient descent algorithm. Since the gradient descent is susceptible to becoming stuck in local optima, an additional model with the hyper parameters of the previous iteration is built and used if the gradient descent cannot provide a model with a smaller root-mean-squared error.

Optimize Infill Criterion: The infill criterion is optimized using a covariance matrix adaption evolution strategy with a population size of 50 (Hansen and Ostermeier 1996). The best found solution from this algorithm is then evaluated with a full TraffSim simulation.

5. RESULTS

In Figure 2 six convergence curves corresponding to the five infill criteria listed in Table 2 are plotted. (Minimal Distance was omitted due to visibility reasons). Two conclusions can be drawn from this figure: Firstly, the difference between most of the infill criteria is different to evaluate and could be influenced by random effects introduced both by the heuristic nature of the optimization and the stochasticity of the simulation. Secondly, several solutions with considerably worse-than-average-quality exist, the frequency of which is illustrated in a histogram (Figure 3) of 1000 uniformly randomly sampled Parameter Vectors. Here each vector was evaluated 20 times, as to remove the impact of simulation specific randomness.

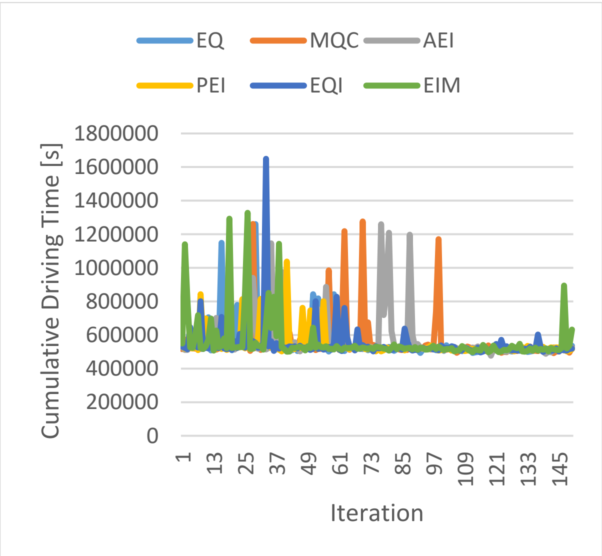

Figure 2: Quality curves of different infill criteria.

The impact of these outliers is twofold. Firstly, they are wasted evaluations that are not only bad in terms of quality but also in terms of execution time since the simulation takes longer to evaluate. Secondly, the strong difference in quality causes undesirable effects in the Gaussian process regression, as the model not only assumes an unjustifiably large variance but also leads the algorithm very closely to the edge of such outliers. Figure 4 shows a Gaussian process regression approximating a one dimensional function with a discrete jump in the fitness landscape. The regression model cannot handle the leap in fitness and oscillates in the vicinity of the jump. If the oscillation is larger than the potential fitness gains of the lower plateau, the algorithms can no longer rely on the regression model and at best fall back into random search and at worst are purposely misled by the model towards the fitness jump.

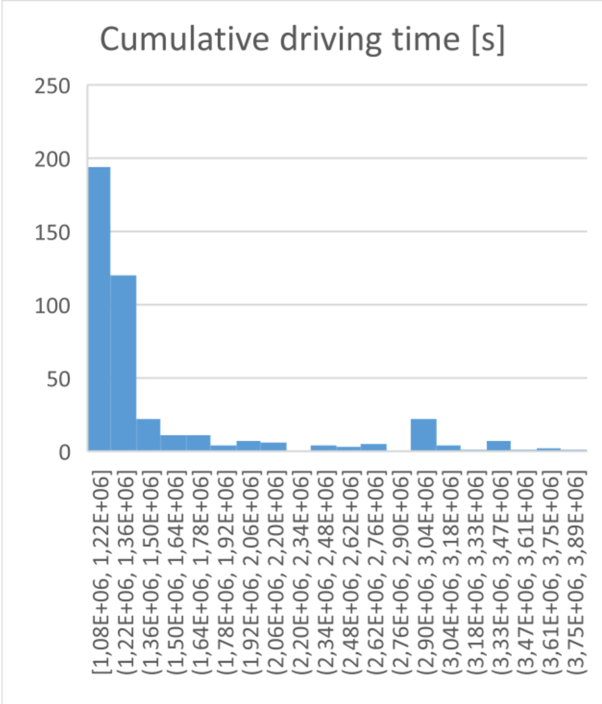

Figure 3: Histogram of cumulative driving times [s]

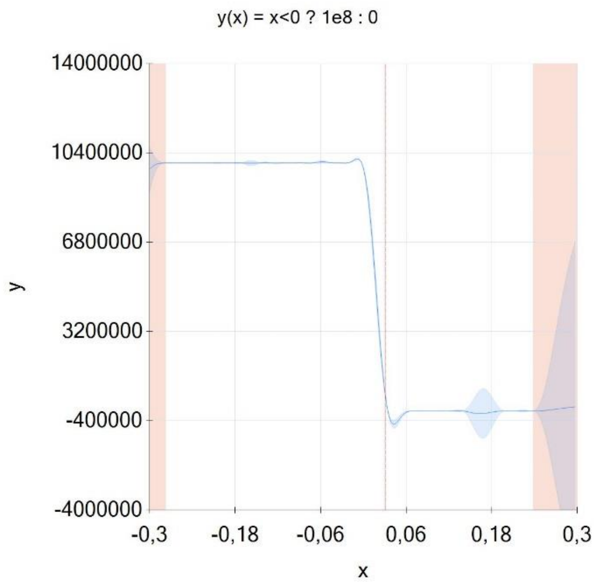

Figure 4: Gaussian process model of a fitness jump

It is therefore imperative to somehow handle these outliers. One option could be to log-transform the fitness landscape, however doing so would create asymmetrical (biased) errors and violate the assumptions of not only the Gaussian process but most regression models in

general. Further, the outliers cannot be attributed to random effects of the traffic simulation and are not distributed randomly in the search space.
In order to explain these outliers Figure 5 presents a projection of the reduced parameter space that was created via t-distributed stochastic neighborhood embedding *tSNE* (Maaten and Hinton 2008). (The most influential parameters were determined by symbolic regression and their impact values were used to weigh the Euclidean distance between points which were fed to the tSNE to determine neighborhood structures). As can be seen the outliers (orange large dots) are located at the edges of the search space. Increasing the lower bounds for the most relevant parameters (rerouting threshold distance (reactive), congestion threshold, congestion level and speed average congestion level) to (1.5, 0.535, 3 and 0.2) respectively removed the regions with extremely bad quality results.

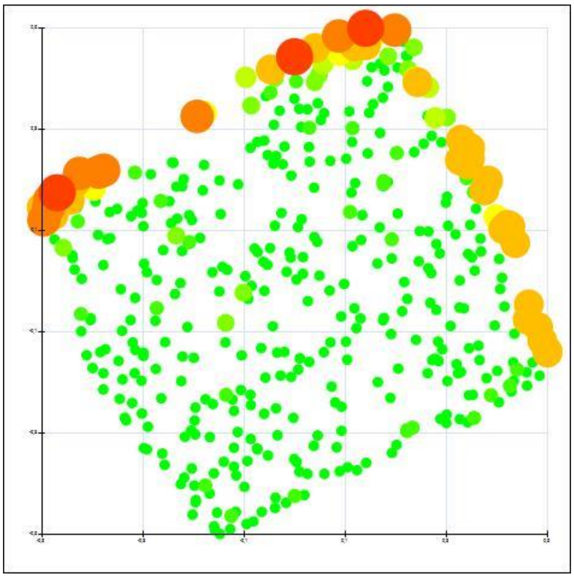

Figure 5: Weighted projection of random samples in the search space. Large red dots denote high objective values while small green denote small objective values

After applying the optimization algorithms to the reduced search space the convergence curves like the one seen in Figure 6 appear to be more similar to what one would expect from an optimization process on a stochastic minimization problem. There is a certain fluctuation that can be attributed to stochastic noise but also a drop in the objective values starting at the 50-samples-mark where the algorithm switches from random sampling to optimizing its infill criterion.
Figure 7 displays the box plots of the best qualities found by optimization algorithm with different infill criteria. Each algorithm configuration was run three times to offset the stochasticity introduced by both the simulation and the heuristic optimization. The Plugin Expected Improvement and the Augmented Expected Improvement appear to be the better infill criteria while the baseline criteria ExpectedQuality and NeighbourDistance are amongst the weaker. It is however important to notice that all achieved quality values lie roughly within the range of one standard deviation of the error measured for each specific problem instance (Standard deviation for the Linz road network with 1500 vehicles: 57629.78s; Standard deviation for the random road network with 1500 vehicles : 25041.30s)

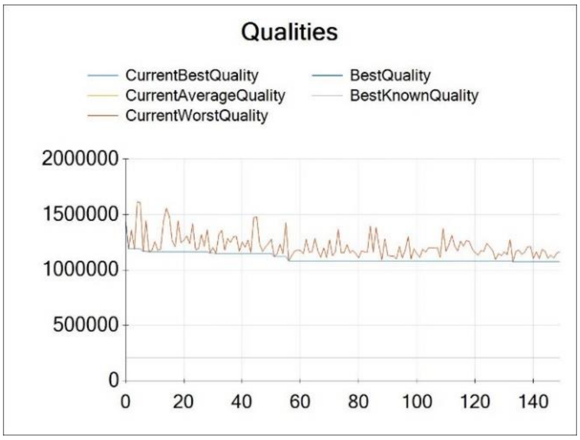

Figure 6: Convergence curve of EGO with Plugin Expected Improvement

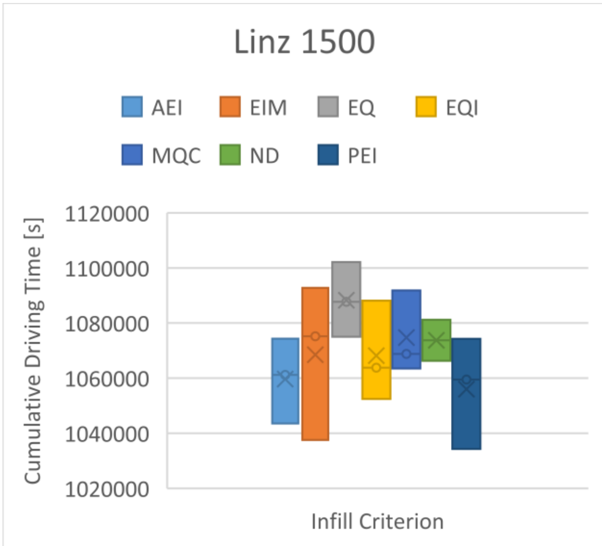

Figure 7: Achieved objective values on Linz with 1500 vehicles (reduced search space)

To gain insight on the almost equal performance of the different infill criteria all evaluated parameter settings from the comparison experiment were used in a linear projection of the search space which is displayed in Figure 8. As can be seen the reduced search space covers a relatively simple part of the fitness landscape where most of the best solutions lie in a valley that stretches diagonally across the projected space.
To gauge the impact of the specific distribution, the uniform distribution of start times was replaced with a normal distribution as described in Section 4. Figure 9 shows a mostly similar fitness landscape for the normal distribution as it was for the uniform distribution. The valley has not changed its position only the area in the lower left corner of the space has gone down in objective

ISBN 978-88-97999-85-0; Affenzeller, Bruzzone, Jiménez, Longo and Piera Eds.

values. Indicating that at least in this specific scenario peaks of high traffic might even be more forgiving to bad parametrization of the routing algorithm than continuous streams of traffic.

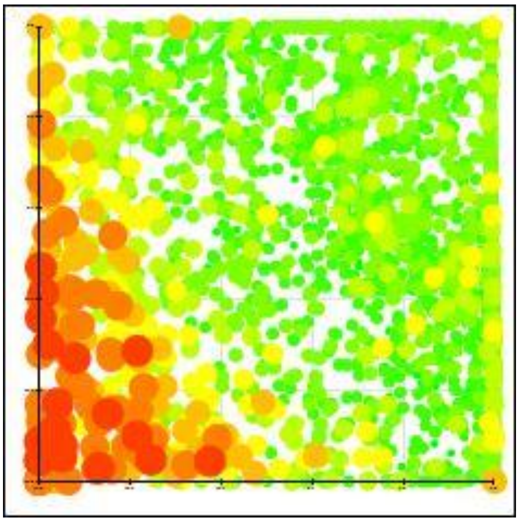

Figure 8: Linear projection of the reduced search space (Linz, 1500 vehicles)

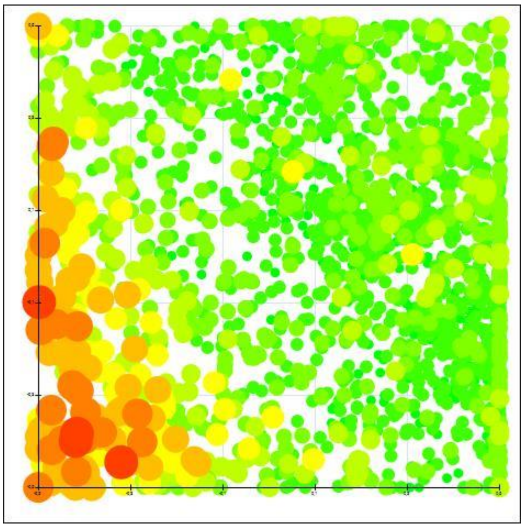

Figure 9: Reduced search space (Linz network, 1500 vehicles, normal distributed start times)

While all infill criteria produce fairly similar objective values, they all are in the lower 10% of the produced points. When looking at the best solution candidates produced by the optimization they all can be found in the valley region of the search space. Therefore the argument can be made that no infill criteria stands out in the boxplots, because the fitness landscape is simple enough that all variants of the EGO algorithm manage to converge towards the favorable region and the ability of an algorithm to converge very precisely is of less significance due to the noise in the fitness function.
Figure 10 uses the same weights as Figure 8 but is created using the samples evaluated while applying the different algorithm configurations to the randomly generated road network. While the same valley structure can be seen, it is important to notice that the valley itself moved a bit, while the red "hill" of unfavorable parameter settings encompasses more area. Some parameter configurations for PCMA* that were optimal for the road map of Linz are now amongst the worse solutions of the search space. Additionally the valley is no longer as well defined as before which can be explained by measuring the different level of noise on both scenarios. Figure 11 shows the best objective values found by all EGO-variants. Again no clear ranking of performance can be established, except for the neighbor distance criterion which seem to perform worst, potentially indicating that more exploitation than in the other scenario is required to find the bottom of the valley.

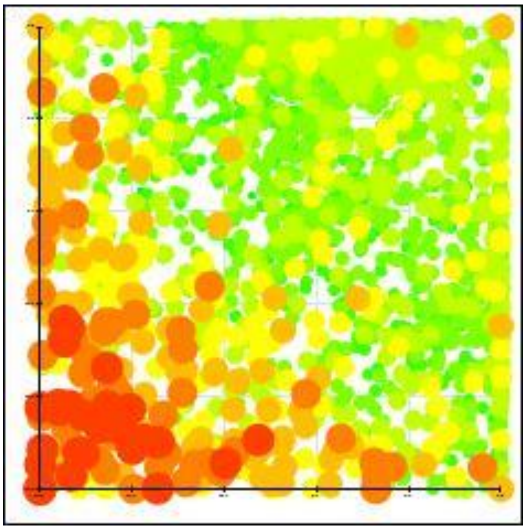

Figure 10: Reduced search space (Random network, 1500 vehicles)

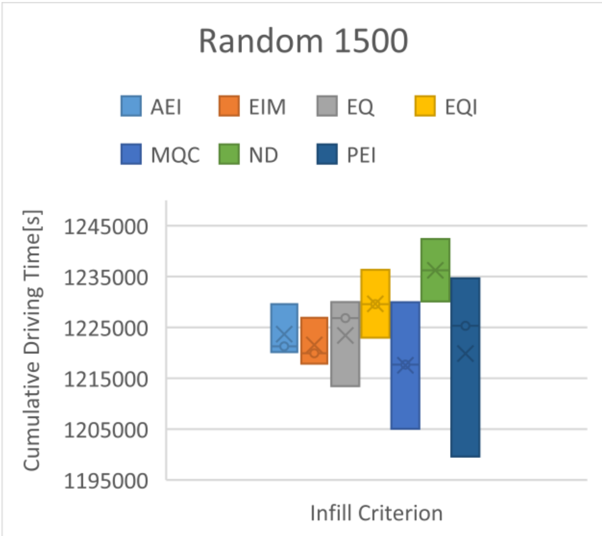

Figure 11: Achieved objective values on the generated road network with 1500 vehicles (reduced search space)

In order to measure the impact of vehicle density on the optimal parameter selection of the routing algorithm the number of vehicles were varied. Figure 12 compares the results achieved on the Linz network for 500 vehicles.

Figure 12: Achieved objective values on Linz with 500 vehicles

As expected, the variance in objective values diminishes with lower numbers of vehicles on the network. The differences in quality by the different infill criteria is neigh irrelevant for scenarios with few vehicles as not as much rerouting needs to happen as can be seen in Table 3 which shows the median cumulative driving times (CDT) for random parameters selections, the best found selection and the CDT achieved without any routing. The column "range" contains the range of all best CDTs achieved by each algorithm run, illustrating that while no essential difference between the performance criteria emerges, all EGO-variants managed to achieve substantial improvement, which increases with the number of vehicles on the network. Interesting to note is that for 500 vehicles on the Linz scenario several parameter settings are even disadvantageous so that the median CDT for random parametrization is effectively worse than disabling the routing algorithm altogether.

Table 3: Comparison of different achieved cumulative driving times

Scenario	Median CDT [s]	Best CDT[s]	Unrouted CDT[s]	Range [s]
Linz 1500	1294988	1034292	3426733	67847
Linz 1000	511211	481152	1243908	16322
Linz 500	220712	206649	208502	927
Random 1500	129967	1199615	4959762	42771

6. CONCLUSIONS

Several conclusions could be drawn from the results presented in the previous section.

- Firstly, that surrogate-assisted algorithms can suffer from unhandled outliers or sharp cliffs in the fitness landscape depending on the type of models that are used.
- Secondly that when it comes to parameter tuning for algorithms, it is rare that all parameters have the same level of impact on the overall performance. It might be a promising extension of existing surrogate-assisted optimization algorithms to include some form of dynamic feature selection.
- Thirdly, an argument can be made that well designed algorithms, be that routing algorithms or otherwise, should have at least somewhat predictable trends in their parameter settings and if the fitness landscape created by these parameters were riddled with local optima, tuning these parameter by hand would be neigh impossible.
- Fourthly it should be noted, that visual inspection of this landscape based on feature selection and dimensionality reduction techniques like principal component analysis, neighborhood component analysis, tSNE or similar methods, can provide significant insight concerning the performance of surrogate-assisted optimization algorithms and the difficulties they face on different problems.
- Fifthly, that when the optimization problem is very noisy, finding an algorithm that outperforms existing ones by a few percent might be fairly inconsequential as these improvements could be offset by small changes to the fitness landscape by choosing a different range or scaling for an input parameter.

The correct selection of parameters for PCMA* is more relevant and more difficult, the more vehicles are on a road network. The impact of the different road networks on the optimal parameter setting is difficult to estimate, but the results somewhat indicate that larger networks might increase the difficulty of selecting the best parameter settings as noise and uncertainty rise.
Future research in this area can be extended in different ways. The number of different road networks tested was fairly limited and larger benchmark instances of road networks and vehicle schedules are required to effectively compare different routing algorithms and their parametrizations. The application of predictive and reactive routing algorithms is geared towards fast decisions that should require considerably less computation time than dynamic traffic assignment, and while surrogate-assisted optimization of the parameters allegedly needs to happen only once per road network, a performance boost and additional insight might be gained from learning the dependence of optimal parameter settings on different geographical properties of a road network.

ACKNOWLEDGMENTS
This work was supported by the European Union through the European Regional Development Fund (EFRE); further information on IWB/EFRE is available at www.efre.gv.at)

REFERENCES
Backfrieder, C., Ostermayer, G., Mecklenbräuker, C. F. (2017). Increased Traffic Flow Through Node-Based Bottleneck Prediction and V2X Communication. IEEE Transactions on Intelligent Transportation Systems, 18(2), 349-363
Backfrieder, C., Mecklenbrauker, C. F., Ostermayer, G. (2013, November). TraffSim--A Traffic Simulator for Investigating Benefits Ensuing from Intelligent Traffic Management. European Modelling Symposium (EMS), 2013 (pp. 451-456). IEEE.
Beyer, H. G., Schwefel, H. P. (2002). Evolution strategies–A comprehensive introduction. Natural computing, 1(1), 3-52.
Boukouvala, F., Ierapetritou, M. G. (2013). Surrogate-based optimization of expensive flowsheet modeling for continuous pharmaceutical manufacturing. Journal of Pharmaceutical Innovation, 8(2), 131-145.
Cox, D. D., John, S. (1992, October). A statistical method for global optimization. Systems, Man and Cybernetics, 1992, IEEE International Conference on (pp. 1241-1246).
Forrester, A. I., Keane, A. J. (2009). Recent advances in surrogate-based optimization. Progress in Aerospace Sciences, 45(1), 50-79.
Greenshields, B. D., Channing, W., Miller, H. (1935). A study of traffic capacity. In Highway research board proceedings (Vol. 1935). National Research Council (USA), Highway Research Board.
Han, Z., Zhang, K., Song, W., Liu, J. (2013). Surrogate-based aerodynamic shape optimization with application to wind turbine airfoils. 51st AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition (p. 1108).
Hansen, N., Ostermeier, A. (1996, May). Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. Evolutionary Computation, 1996., Proceedings of IEEE International Conference on (pp. 312-317). IEEE.
He, X. (2014). Simulation-based optimization of transportation systems: Theory, surrogate models, and applications Doctoral dissertation. University of Maryland
Huang, D., Allen, T. T., Notz, W. I., Zeng, N. (2006). Global optimization of stochastic black-box systems via sequential kriging meta-models. Journal of global optimization, 34(3), 441-466.
Kesting, A., Treiber, M., Helbing, D. (2010). Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 368(1928), 4585-4605.
Loshchilov, I., Schoenauer, M., Sebag, M. (2013, July). Intensive surrogate model exploitation in self-adaptive surrogate-assisted cma-es (saacm-es). Proceedings of the 15th annual conference on Genetic and evolutionary computation (pp. 439-446). ACM.
Janson, B. N. (1991). Dynamic traffic assignment for urban road networks. Transportation Research Part B: Methodological, 25(2), 143-161.
Johnson, M. E., Moore, L. M., Ylvisaker, D. (1990). Minimax and maximin distance designs. Journal of statistical planning and inference, 26(2), 131-148.
Jones, D. R., Schonlau, M., Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. Journal of Global optimization, 13(4), 455-492.
Lindorfer, M., Backfrieder, C., Kieslich, C., Krosche, J., Ostermayer, G. (2013, November). Environmental-Sensitive Generation of Street Networks for Traffic Simulations. European Modelling Symposium (EMS), 2013 (pp. 457-462). IEEE.
Maaten, L. V. D., Hinton, G. (2008). Visualizing data using t-SNE. Journal of Machine Learning Research, 9(Nov), 2579-2605.
Osborne, M. A., Garnett, R., Roberts, S. J. (2009, January). Gaussian processes for global optimization. 3rd international conference on learning and intelligent optimization (LION3), 1-15.
Picheny, V., Ginsbourger, D., Richet, Y. (2010). Noisy expected improvement and on-line computation time allocation for the optimization of simulators with tunable fidelity.
Picheny, V., Wagner, T., Ginsbourger, D. (2013). A benchmark of kriging-based infill criteria for noisy optimization. Structural and Multidisciplinary Optimization, 48(3), 607-626.
Schrank, D., Eisele, B., Lomax, T. (2012). TTI's 2012 urban mobility report. Texas A&M Transportation Institute. The Texas A&M University System.
Sun, C., Jin, Y., Zeng, J., Yu, Y. (2015). A two-layer surrogate-assisted particle swarm optimization algorithm. Soft computing, 19(6), 1461-1475.
Wang, G. G., Shan, S. (2007). Review of metamodeling techniques in support of engineering design optimization. Journal of Mechanical design, 129(4), 370-380.