



**Technical report for  
Collaborative Research Center  
SFB 876**

**Providing Information by Resource-  
Constrained Data Analysis**

December 2018

Part of the work on this report has been supported by Deutsche Forschungsgemeinschaft (DFG) within the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Analysis".

Speaker: Prof. Dr. Katharina Morik  
Address: Technische Universität Dortmund  
Fachbereich Informatik  
Lehrstuhl für Künstliche Intelligenz, LS VIII  
D-44221 Dortmund

# Contents

<b>1</b>	<b>Subproject A1</b>	<b>1</b>
1.1	Sebastian Buschjäger . . . . .	2
1.2	Alexander Lochmann . . . . .	6
1.3	Lukas Pfahler . . . . .	10
<b>2</b>	<b>Subproject A2</b>	<b>15</b>
2.1	Amer Krivošija . . . . .	16
<b>3</b>	<b>Subproject A3</b>	<b>21</b>
3.1	Andrea Bommert . . . . .	22
3.2	Jakob Richter . . . . .	26
<b>4</b>	<b>Subproject A4</b>	<b>31</b>
4.1	Stefan Böcker . . . . .	32
4.2	Markus Buschhoff . . . . .	36
4.3	Robert Falkenberg . . . . .	40
4.4	Pascal Jörke . . . . .	44
4.5	Mojtaba Masoudinejad . . . . .	48
4.6	Karthik Aswin Venkatapathy Ramachandran . . . . .	52
4.7	Janis Tiemann . . . . .	56
<b>5</b>	<b>Subproject A6</b>	<b>61</b>
5.1	Andre Droschinsky . . . . .	62
5.2	Matthias Fey . . . . .	66
5.3	Christopher Morris . . . . .	70
<b>6</b>	<b>Subproject B1</b>	<b>75</b>
6.1	Salome Horsch . . . . .	76

<b>7</b>	<b>Subproject B2</b>	<b>81</b>
7.1	Kuan-Hsun Chen . . . . .	82
<b>8</b>	<b>Subproject B3</b>	<b>87</b>
8.1	Amal Saadallah . . . . .	88
8.2	Jacqueline Schmitt . . . . .	92
8.3	Mario Wiegand . . . . .	96
<b>9</b>	<b>Subproject B4</b>	<b>101</b>
9.1	Karsten Heimann . . . . .	102
9.2	Florian Liedmann . . . . .	106
9.3	Benjamin Sliwa . . . . .	110
9.4	Tim Vranken . . . . .	114
<b>10</b>	<b>Subproject C1</b>	<b>119</b>
10.1	Sibylle Hess . . . . .	120
10.2	Marc Schulte . . . . .	124
10.3	Henning Timm . . . . .	128
<b>11</b>	<b>Subproject C3</b>	<b>133</b>
11.1	Kai Brügge . . . . .	134
11.2	Mirko Bunse . . . . .	138
11.3	Mirco Hünnefeld . . . . .	142
11.4	Maximilian Meier . . . . .	146
11.5	Maximilian Nöthe . . . . .	150
11.6	Kevin Schmidt . . . . .	154
11.7	Jan Soedingrekso . . . . .	158
<b>12</b>	<b>Subproject C5</b>	<b>163</b>
12.1	Thomas Lindemann . . . . .	164
12.2	Margarete Schellenberg . . . . .	168

12.3 Holger Stevens . . . . .	172
-------------------------------	-----





Subproject A1  
Data Mining for Ubiquitous System Software

Katharina Morik      Olaf Spinczyk

# Machine Learning on Embedded Systems

Sebastian Buschjäger  
Artificial Intelligence Group, Chair 8  
Technical University Dortmund  
sebastian.buschjaeger@tu-dortmund.de

With increasing volumes in data and more sophisticated machine learning algorithms, the demand for fast and energy efficient computation systems is also growing. To meet this demand, two approaches are possible: First, machine learning algorithms can be tailored specifically for the hardware at hand. Second, instead of changing the algorithm we can change the hardware to suit the machine learning algorithms better. This report briefly discusses my last years' work which focused largely on the first approach and quickly outlines some ideas for future research.

## 1 Introduction

To make machine learning universally applicable, we need to bring its algorithms to small and embedded devices including both - the training and the application of models. From a computer architectural point of view, we may optimize these two aspects separately. In model application we rapidly apply an already trained model for predictions and thus focus on the optimization of inference. In model training however, we would like to train models on small devices directly, so that these devices dynamically adjust their prediction rules for new data.

## 2 Machine learning for Embedded Devices

In the previous year we looked at filtering of sensor data by using Random Forests for data driven filtering rules. Our goal is to apply a given Random Forest model continuously on a small devices, to pre-filter sensor measurements before transmitting them to a central server. Random Forests are among the most widely used Machine Learning models and offer compelling theoretical properties. A Random Forest consist of  $M$  Decision Trees, which are trained on different subsets of features and/or different subsets of the data. A majority vote is used to combine the individual predictions into a single one. Integral to the implementation of Random Forests is the implementation of each single tree. A tree consist of a series of binary decisions of the form ' $x_i < t$ ', where  $i$  is a specific feature of the input  $\vec{x} \in \mathbb{R}^d$  and  $t \in \mathbb{R}$  is a threshold. Depending on the outcome of that decision, we move forward and look at the left or right child of that node. We repeat this recursively until we hit a leaf-node, which then contains the final prediction.

In [3] we discussed four different ways on how to implement Decision Trees, taking data cache, instruction cache and vectorization units into account. This year, we built onto that work by considering the caching behaviour of modern CPUs more explicitly. In cooperation with the embedded systems groups (chair 12) we derived two algorithms which optimize the memory layout of Decision Trees with respect to the instruction and data cache in [2]. In an extensive set of experiments, we were able to show that our method improves the evaluation speed of Decision Trees by a factor of 2 – 6 compared to a reference implementation and by a factor around 1500 compared to `sklearn` in the real-time setting. Moreover, the set of experiments showed that the optimal memory layout is dependent on the CPU architecture and the specific CPU model. Thus, we implemented a code-generator which generates optimized code for a given architecture and specific CPU parameters, such as cache size<sup>1</sup>.

## 3 Machine learning on Embedded Devices

Going from model application to model training, I stayed in the realm of embedded systems and ensemble methods. In their most general form, ensemble methods produce a predictive model  $f(x)$  by combining several base models in a weighted, linear fashion

$$f(x) = \sum_{i=1}^T \lambda^i h^i(x)$$

where  $h^i \in \mathcal{H}$  is a set of base learners and  $\lambda^i \in \mathbb{R}$  are the weights of each base model. For example, a Random Forest model uses Decision Trees as base learners which are

---

<sup>1</sup>Our software is available at <https://bitbucket.org/sbuschjaeger/arch-forest/src>

trained on different subsets of features and/or the data, where each tree is weighted equally with  $\lambda^i = 1/T$ . At first it seems counter-intuitive to use ensemble methods on small devices, because  $T$  instead of one model needs to be trained. However, the general analysis of ensemble methods allow for any set of base learners  $\mathcal{H}$ . When considering small devices, we can restrict ourselves specifically to a set of base learners with e.g. only integer parameters or fixed memory requirements. This way, we can use a rather weak, but efficiently trainable set of base learners and transform it into a strong learner using ensemble techniques. This gives us two parameters which can be used to fine-tune our ensemble to the specific device at hand, namely the set of weak learners and the size of the ensemble.

In order to bring ensemble methods to small devices, I noticed two major challenges: First, there is a lack of software tools and implementations which can actually run on small devices. To the best of my knowledge, all major software tools such as RapidMiner, Weka or XGBoost focus on Desktop PCs / Servers (RapidMiner, Weka) or GPU acceleration (XGBoost). A notable exception in this context is Microsoft's EdgeML<sup>2</sup>, which offers specialized versions of K-NN and Decision Trees for embedded systems. However, training of these models is still performed on a larger PC, but only its application is meant to be done on a small system. Second, there exists a vast landscape of different embedded devices and architectures, which often require to fine-tune the implementation to the specifics of each system. To cope with these challenges, I decided to implement my own software library called `ensembles`<sup>3</sup> which specifically targets embedded systems. To this date (December 2018), this library contains multiple base-learners, namely Decision Stump, Decision Trees, Linear Regression and Gaussian Processes, as well as Boosting and Bagging. Most methods are available as a (classic) batch learning algorithm for reference and an online version which can run on an embedded device. I installed an automatic pipeline, to test and cross-compile the library for Intel x86, ARMv6 and ARMv7 CPUs. The next step is to include more esoteric CPU architectures into this set, such as the ESP8266<sup>4</sup> and to include automatic measuring of the power consumption.

## 4 Future research

In contrast to the mostly practical work in the second half of 2018, I also investigated the theoretical foundations of ensemble methods. More specifically, I looked at the coordinate descent view of boosting, in which we view each base learner  $h \in \mathcal{H}$  as a coordinate of  $f$ . Subsequently, the computation of the respective weights  $\lambda$  can be seen as a coordinate descent algorithm. This view allows for greater freedom in the design of boosting algorithm, especially in the context of limited memory and limited

---

<sup>2</sup><https://github.com/Microsoft/EdgeML>

<sup>3</sup><https://bitbucket.org/sbuschjaeger/ensembles/src>

<sup>4</sup><https://de.wikipedia.org/wiki/NodeMCU>

computational resources. More specifically, I want to use this view to analyse stochastic boosting algorithms [1, 6] which only consider a fraction of the training data, as well as distributed boosting algorithms [4, 5, 7] which train multiple base learner in parallel. A paper on this subject is currently in progress. Interestingly, the coordinate descent view of boosting also encapsulates Decision Trees. In other words, DT induction algorithms such as ID3 or C4.5 can also be viewed as boosting algorithms, which may enable new generalization bounds for DTs.

As a second research direction, I want to extend my research in the area of model application on small devices. In the context of Random Forests, we are currently investigating the possibility to reduce their memory requirements by removing duplicate sub-trees from a forest. This problem can be either viewed as a hard matching problem, where two sub-trees are the same if they use the same combination of features and thresholds, or as a fuzzy matching problem, in which two sub-trees are equal if they have the same prediction on the same input.

## References

- [1] Joseph K Bradley and R Schapire. Filterboost: Regression and classification on large datasets. *Advances in neural information processing systems*, 20:185–192, 2008.
- [2] Sebastian Buschjäger, Kuan-hsun Chen, Jian-jia Chen, and Katharina Morik. Realization of Random Forest for Real-Time Evaluation through Tree Framing. *The IEEE International Conference on Data Mining series (ICDM)*, 2018.
- [3] S. Buschjäger and K. Morik. Decision tree and random forest implementations for fast filtering of sensor data. *IEEE Transactions on Circuits and Systems I: Regular Papers*, PP(99):1–14, 2017.
- [4] Jeff Cooper and Lev Reyzin. Improved algorithms for distributed boosting. In *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017*, volume 2018-Janua, pages 806–813, 2018.
- [5] Wei Fan, Salvatore J Stolfo, and Junxin Zhang. The application of AdaBoost for distributed, scalable and on-line learning. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '99*, number January 1999, pages 362–366, 1999.
- [6] Jerome H. Friedman. Stochastic gradient boosting. *Computational Statistics and Data Analysis*, 2002.
- [7] Aleksandar Lazarevic and Zoran Obradovic. The distributed boosting algorithm. *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '01*, pages 311–316, 2001.

# Optimized Bug Reports for easy Detection of Locking Bugs using LockDoc

Alexander Lochmann  
Arbeitsgruppe Eingebettete Systemsoftware  
Technische Universität Dortmund  
alexander.lochmann@tu-dortmund.de

LockDoc [3, 4] is a trace-based approach to automatically derive locking rules for data types in operating systems. One outcome of LockDoc are reports of potential bugs – one for each observed data type. Applying LockDoc on Linux revealed that those reports could be long and not easy to assess [3,4]. This report presents a new method of compressing the reports, and giving the developer a better overview over the data.

## 1 Introduction

LockDoc [3, 4] generates reports of potential bugs, the so-called counterexamples, as HTML websites. These makes them easily exchangable with developers while using modern web techniques to present those counterexamples. The reports generated by LockDoc fullfill three tasks: a) Group the results by data type, element, and access type, i.e. r or w, b) display the winning locking hypothesis, c) present the stack traces that lead to the memory access, and d) list the locks actually held. Task c) and d) are the main objectives since they allow an easy assessment by the developer. The first two tasks are accomplished by the heading shown in Figure 1. It clearly separates the counterexamples as well as presenting the winning hypothesis. The tasks c) and d) are not well covert: The first column in Figure 1 shows the stack trace leading to the memory accesses. Whereas, columns three and four show how often which set of locks that were actually held. That table might become confusing for the reader if there are

Hypothesis 2: When writing <code>inode.i_blocks</code> the following locks should be held: <code>EMBSAME(inode.i_lock[w])</code> 93.56% (189 out of 202 mem accesses under Locks)			
Stacktrace (00 = stackframe where access occurred)	ID	Occurrences	Locks actually held (in order locks were taken)
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.1	10	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.2	1	01: <a href="#">EMBOTHER(i_rwsem[w])</a> 02: <a href="#">EMBSAME(i_rwsem[w])</a> 03: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.3	1	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">EMBOTHER(i_rwsem[w])</a> 03: <a href="#">uevent_sock_mutex:163(mutex[w])</a> 04: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.4	1	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">EMBOTHER(i_rwsem[w])</a> 03: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.5	10	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.6	1	01: <a href="#">EMBOTHER(i_rwsem[w])</a> 02: <a href="#">EMBSAME(i_rwsem[w])</a> 03: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.7	1	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">EMBOTHER(i_rwsem[w])</a> 03: <a href="#">uevent_sock_mutex:163(mutex[w])</a> 04: <a href="#">41(hardirq[w])</a>
13: <a href="#">entry_SYSENTER_32:458</a> 12: <a href="#">do_fast_syscall_32:384</a> 11: <a href="#">do_syscall_32_irqs_on:322</a> 10: <a href="#">Sys_write:599</a> 09: <a href="#">SYSC_write:607</a> 08: <a href="#">vfs_write:560</a> 07: <a href="#">_vfs_write:512</a> 06: <a href="#">new_sync_write:499</a> 05: <a href="#">generic_file_write_iter:2979</a> 04: <a href="#">_generic_file_write_iter:2951</a> 03: <a href="#">generic_perform_write:2826</a> 02: <a href="#">shmem_write_begin:2211</a> 01: <a href="#">shmem_getpage:123</a> 00: <a href="#">shmem_getpage_gfp:1790</a>	2.8	1	01: <a href="#">EMBSAME(i_rwsem[w])</a> 02: <a href="#">EMBOTHER(i_rwsem[w])</a> 03: <a href="#">41(hardirq[w])</a>

Figure 1: An example of a Bug Report for writing `inode.i_blocks`. Column *Stacktrace* shows the code location where a suspicious memory access occurred. Columns three and four show how often which set of locks were actually held.

a large number of stack traces, or a large number of lock sets. In that case, the reader has to navigate back and forth to compare the results with each other, and to look up the winning hypothesis again. Moreover, this way of presenting the counterexamples lacks the relations between the different stack traces. The stack traces can be almost identical except for the last element, for example. Thus, I present two more concise ways of presenting counterexamples in the following sections Section 2 and Section 3. Section 4 explains a way of combining both ideas to further improve the presentation.

## 2 Tree-based Visualization

The tree-based visualization combines both the location of a bug and the locks actually held in one figure. The *Treant* JavaScript library is used to draw the trees [1]. A tree gives a developer an easy and quick overview of the counterexamples for one particular element of a data type, and where they occurred. It furthermore displays the relations between the counterexamples.

As can be seen in the example in Figure 2b each path from the root to a leaf represents one stack trace. Each node corresponds to one element of the trace, i.e. a code location. A leaf represents the code location, i.e. file and line, where suspicious memory accesses happened. A leaf also contains a set of counterexamples. This includes the

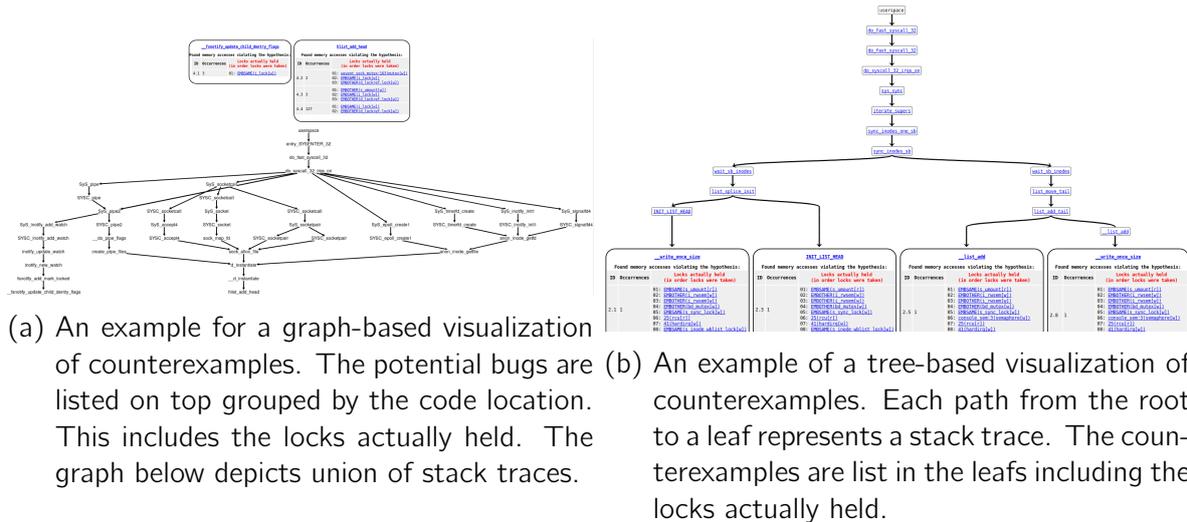


Figure 2: Examples of different kind of visualizations used to present counterexamples.

aforementioned set of lock combinations of actually held locks as well as the number of occurrences. However, a tree might become confusing if a) many disjunct stack traces exist and/or b) many deep stack traces exist. In that case, the tree tends to expand both vertically and horizontally. This makes it hard for a developer to assess the counterexamples. He or she has to scroll to reach the leaves. The growth is due to the fact that the tree-based visualization may contain leaves with redundant code locations. The nature of trees favors the existence of redundant leaves: Paths may branch but are not allowed to be merged again. Hence, if two stack traces have the same prefix and the same suffix but differ in between they are presented as two paths branching after the prefix.

To sum it up, the tree-based presentation allows for an easy assessment of the counterexamples, and the relations among them. However, a growing number of unique stack traces makes this approach impractical since the trees rapidly grow in size.

### 3 Graph-based Visualization

The graph-based presentation in contrast uses a graph to depict the stack traces as shown in Figure 2a. The graph is plotted using *Cytoscape JS*<sup>1</sup>, a JavaScript version of the graph library Cytoscape [2]. A graph, by definition, allows to join paths again which elides redundant nodes. This makes the visualization more compact when the amount of stack traces increases. In contrast to the tree view, it is not guaranteed that all leaves of the graph are arranged at the bottom. The list of counterexamples might be drawn somewhere in the graph. To overcome that issue, the lists of counterexamples are

<sup>1</sup>Cytoscape.js, <http://js.cytoscape.org/> (Accessed January 8th, 2019)

therefore placed above the call graph shown Figure 2a. The counterexamples are grouped by the code location, i.e. function and line, that performed the suspicious access. This way the developer gets a quick overview of the variety of the counterexamples without scrolling to the bottom. However, the visualization lacks the relations between counterexamples and stack traces.

## 4 Visualization in Practice

Both ways of presenting counterexamples have their advantages and disadvantages. None is favorable over the other. Therefore, a combined approach is used: LockDoc chooses one visualization per tuple of (data type element,  $\{r,w\}$ ). If the amount of leafs in the tree-based visualization exceeds a threshold, the graph-based one is chosen. Experiments showed that trees with 8 or more leafs are better presented as graphs.

## References

- [1] Treant.js - javascript library for drawing tree diagrams, 2018 (accessed January 8th, 2018).
- [2] Max Franz, Christian T. Lopes, Gerardo Huck, Yue Dong, Onur Sumer, and Gary D. Bader. Cytoscape.js: a graph theory library for visualisation and analysis. *Bioinformatics*, 32(2):309–311, 2016.
- [3] Alexander Lochmann, Horst Schirmeier, Hendrik Borghorst, and Olaf Spinczyk. LockDoc: Trace-based analysis of locking in the linux kernel. In *Proceedings of the 14th ACM SIGOPS/EuroSys European Conference on Computer Systems (EuroSys '19)*, New York, NY, USA, 2019. ACM Press. to appear.
- [4] Katharina Morik and Wolfgang (Editors) Rhode. Technical report for collaborative research center sfb 876 - graduate school. Technical Report 2, TU Dortmund University, December 2017. Section: Subproject A1 - Alexander Lochmann.

# The Spectrum of Eigenvalues in Machine Learning Models

Lukas Pfahler  
Artificial Intelligence Group, LS 8  
TU Dortmund University  
lukas.pfahler@tu-dortmund.de

## 1 Introduction

The underlying theme of my research into machine learning methods seems to revolve around eigenvalues. The behavior of machine learning models is often governed by the spectrum of eigenvalues or singular values of its parameter matrices or its underlying training data matrices. This offers possibilities for bounding the generalization abilities of these models [2], introducing regularization [7] or new learning approaches [6]. In the following we will see work in the field of kernel methods, where the focus is on the eigenvalues of the kernel matrix, as well as work on modern neural networks, where we inspect the – possibly very large – weight matrices.

## 2 Kernel Learning

Kernel methods are a popular choice for classification problems, but when solving large-scale learning tasks they quickly become infeasible. The essential component of any kernel method is the kernel matrix. Algorithmically, the biggest challenge is the quadratic complexity of computing and storing the kernel matrix. If the eigenvalues of the kernel matrix decay fast, it is feasible to compute an accurate low-rank approximation of the matrix. The Nyström method that approximates the kernel matrix using only a smaller sample of the kernel matrix has been proposed. We sample a constant number of rows of the full kernel matrix and use these rows to build a low-rank approximation. Other techniques to speed up kernel learning include stochastic first order optimization and

conditioning. We have introduced Nyström-SGD [6], a learning algorithm that trains kernel classifiers by minimizing a convex loss function with conditioned stochastic gradient descent while exploiting the low-rank structure of a Nyström kernel approximation. Our experiments suggest that the Nyström-SGD enables us to rapidly train high-accuracy classifiers for large-scale classification tasks.

For this algorithm, it is essential to compute the exact eigenvalues of the Nyström-approximation. Only this allows us to converge to the empirical risk minimizer in a numerically-stable manner. The special structure of the Nyström kernel matrix allows us to compute the eigenvalues faster than the usual  $\mathcal{O}(n^3)$ .

### 3 Inspecting Singular Values in Large Neural Networks

The behavior of deep networks is governed by their weight matrices  $W$  and activation functions  $\sigma$ , as each layer computes a new representation  $h(x) = \sigma(W h'(x))$ . We are interested in the singular values of the linear operators  $W$ . Instead of considering the singular values of  $W$ , we can equivalently analyze the eigenvalues of  $WW^T$ . Theoretical results show the importance of the singular values for generalization abilities [2], as the state-of-the-art bounds are functions of Schatten-norms, which in turn are functions of the singular values.

We want to visually inspect the eigenvalue distribution of the operators by creating histograms. Computing eigenvalues exactly is a costly operation, particularly because the dimension of hidden layers in state-of-the-art convolutional neural networks often exceeds 100,000. Fortunately, when we are interested in a histogram, we only need to know the number of eigenvalues that fall into the bins of the histogram. We derive a method that does not need to compute all the eigenvalues exactly. It is based on two techniques: For estimating a small number of the largest eigenvalues exactly, we use ARPACK, a truncated eigenvalue decomposition method that does not require the matrix explicitly, but accesses it only via matrix-vector products [4]. For estimating the remaining histogram, we use a method based on matrix Chebyshev approximations and Hutchinson's trace estimator [5]. It also only accesses the matrix via matrix-vector products. Ideally we would like to use the cheaper stochastic estimator for all bins, but additive errors are much worse for the histogram bins with very few eigenvalues. Hence we settle for a compromise where we compute a few bins exactly, exploiting the rapid decay of eigenvalue counts.

The highest-dimensional linear operators in deep networks used in production environments are likely convolution layers. These layers transform feature maps with number of raw-features in the input- and output feature often exceeding 100k. This is feasible because the convolution operator is not implemented as matrix-vector multiplication with dense weight matrices, but specialized and highly-optimized convolution routines are

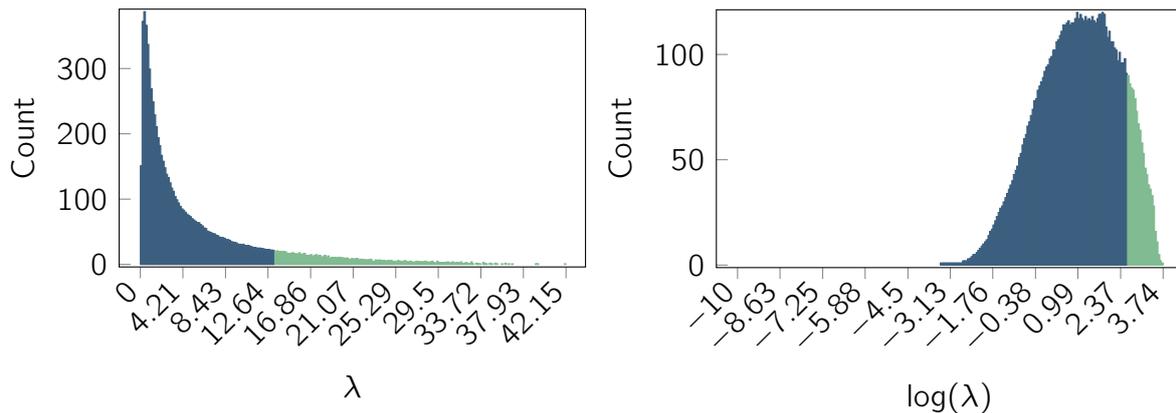


Figure 1: Histogram of a 3x3 convolution linear convolution layer that maps from 8,112 features to 32,448 features after 90 epochs of imagenet training.

used. We can use these same routines when we estimate the eigenvalues of the linear maps of network layers, as we only access the matrix via matrix vector products.

Figure 1 shows preliminary results, where the eigenvalues of a convolution layer in a deep network were analyzed. The green part of the plot shows the eigenvalues that have been computed exactly with the linear algebra library, whereas the blue part shows the eigenvalue counts that have been approximated using the stochastic estimator. The results are based on our own unpublished python library that can analyze the linear operators in any network implemented in pyTorch. The eigenvalues seem to follow a log-normal distribution and their counts decay rapidly. However, the matrix has full rank.

## 4 Future Research

In the future I want to expand the analysis of eigenvalues in deep networks based on stochastic estimators into two directions: regularization and covariance structure.

In order to apply stochastic estimators for **regularization** in models that are trained using stochastic gradient descent, it is essential that the estimators are unbiased and differentiable and that the gradients are unbiased as well. There are extensions to the approach we use for the histograms that provide just that [1, 3]. This allows us to test regularizers of the form  $R(A) = \sum_i f(\lambda_i(A))$  where  $f$  is a real-valued function applied to each eigenvalue  $\lambda_i(A)$  of  $A$ . Overall the computational efficiency of stochastic regularizers should exceed the performance of comparable regularizers that require computing singular value decompositions at each optimization step. However they do not allow more general forms of regularizers, e.g.  $R(A) = \sum_i f_i(\lambda_i(A))$ .

The generalization performance of linear models can be characterized by the decay of eigenvalues of the **covariance** of the data. If there is exponential decay, then it is possible to obtain generalization bounds that show a fast rate of convergence  $\mathcal{O}(1/n)$  in contrast to the usual  $\mathcal{O}(\sqrt{1/n})$ . It is unclear if a similar result can be shown for deep networks. As a first step I want to analyze the covariance structure within a deep network and visualize the eigenvalues of the covariance matrices of intermediate representations. Computing the full covariance matrices is infeasible, but computing unbiased estimators is possible. These unbiased estimators can be used to compute an unbiased estimate of the spectrum [1].

Ultimately a combination of the two directions is plausible: regularize such that the covariance structure in intermediate layers follows nice decay laws.

## References

- [1] Ryan P. Adams, Jeffrey Pennington, Matthew J. Johnson, Jamie Smith, Yaniv Ovdia, Brian Patton, and James Saunderson. Estimating the Spectral Density of Large Implicit Matrices. Technical Report 3, 2018.
- [2] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-Independent Sample Complexity of Neural Networks. In Philippe Bubeck, Sebastien and Perchet, Vianney and Rigollet, editor, *Proceedings of the 31st Conference On Learning Theory*, number 1, pages 297–299. PMLR, 2018.
- [3] Insu Han, Haim Avron, and Jinwoo Shin. Optimizing Spectral Sums using Randomized Chebyshev Expansions. pages 1–21, 2018.
- [4] R B Lehoucq, D C Sorensen, and C Yang. *ARPACK Users' Guide: Solution of Large-scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. Software, Environments, Tools. Society for Industrial and Applied Mathematics, 1998.
- [5] Edoardo Di Napoli, Eric Polizzi, and Yousef Saad. Efficient estimation of eigenvalue counts in an interval. *Numerical Linear Algebra with Applications*, 2016.
- [6] Lukas Pfahler and Katharina Morik. Nyström-SGD : Rapidly Learning Kernel-Classifiers with Conditioned Stochastic Gradient Descent. In *ECML PKDD 2018*, 2018.
- [7] Lukas Pfahler, Katharina Morik, Frederik Elwert, Samira Tabti, and Volkhard Krech. Learning Low-Rank Document Embeddings with Weighted Nuclear Norm Regularization. In *Proceedings of the 4th DSAA*, 2017.





Subproject A2  
Algorithmic aspect of learning methods in  
embedded systems

Christian Sohler      Jens Teubner

# Theoretical and Experimental Bounds on the Fréchet Distance Distortion

Amer Krivošija

Lehrstuhl für Effiziente Algorithmen und Komplexitätstheorie  
Technische Universität Dortmund  
amer.krivosija@tu-dortmund.de

The Fréchet distance is a popular distance measure for curves, but the computation complexity to determine the similarity between two given curves poses considerable computational challenges in practice. We studied distortion of the probabilistic embedding that results from projecting the curves to a randomly chosen line, both from theoretical and experimental view, on a set of realistic input curves.

## Introduction

The Fréchet distance is a distance measure for curves which naturally lends itself to fundamental computational tasks, such as clustering, nearest-neighbor searching, and spherical range searching in the corresponding metric space. However, their inherent complexity poses considerable computational challenges in practice. Indeed, spherical range searching under the Fréchet distance was recently the topic of the yearly ACM SIGSPATIAL GISCUP competition [6], highlighting the relevance and the difficulty of designing efficient data structures for this problem. At the same time, Afshani and Driemel show lower bounds on the space-query-tradeoff in the pointer model [1] that demonstrate that this problem is even harder than simplex-range searching.

The computational complexity of computing a single Fréchet distance between two given curves is a well-studied topic. It is believed that it takes time that is quadratic in the length of the curves and this running time can be achieved by applying dynamic programming. In the literature the case of 1-dimensional curves under the continuous Fréchet distance stands out, as there are no lower bounds known on computing the continuous

Fréchet distance between 1-dimensional curves. It has been observed that the problem has a special structure in this case. Clustering under the Fréchet distance can be done efficiently for 1-dimensional curves [5], but seems to be harder for curves in the plane or higher dimensions. Bringmann and Künnemann use projections to lines to speed up their approximation algorithm for the Fréchet distance [2]. It may be assumed that the problem for the 1-dimensional curves is significantly easier, but in the general case, there are no algorithms known which are faster for 1-dimensional curves than for curves in higher dimensions. In practice, it is very common to separate the coordinates of trajectories to simplify computational tasks. It seems that in practice the inherent character of a trajectory is often largely preserved when restricted to one of the coordinates of the ambient space. Mathematically, this amounts to projecting the trajectory to a line.

## Theoretical bound

Given two polygonal curves  $P$  and  $Q$  with  $t$  vertices each from  $\mathbb{R}^d$ , where  $d \in \{2, 3, 4, 5\}$ . Consider sampling a unit vector  $\mathbf{u}$  in respective  $\mathbb{R}^d$  uniformly at random, and let  $P'$  and  $Q'$  be the projections of the two curves to the line supporting  $\mathbf{u}$ . We observed that Fréchet distance always decreases when the curves are projected to a line, and showed that if the curves  $P$  and  $Q$  are  $k$ -packed for constant  $k$ , then, with constant probability, the discrete Fréchet distance between the curves  $P$  and  $Q$  degrades by at most a linear factor in  $t$ . This ratio we call distortion  $c = d_F(P', Q') / d_F(P, Q)$ . Note that under this definition it is  $c \in [0, 1]$ . The main claim on upper bound is stated by Theorem 1.<sup>1</sup> We also present a matching lower bound on  $c$  [4]. A curve is called  $k$ -packed for a value  $c > 0$  if the length of the intersection of the curve with any ball of any radius  $r$  is at most  $kr$ , and was proved useful as a realistic input assumption [3]. Our study concentrates onto the discrete Fréchet distance, but we expect that our techniques can be extended to the case of the continuous Fréchet distance.

**Theorem 1.** *Given  $k \geq 2$ , for any two polygonal  $k$ -packed curves  $P$  and  $Q$  from  $\mathbb{R}^d$ , where  $d \in \{2, 3, 4, 5\}$ , and for any  $\gamma \in (0, 1)$  it holds that*

$$\Pr[c \leq 1/(k' \cdot t)] \leq \gamma, \quad (1)$$

where the constant  $k'$  is  $(12k + 16)/\gamma$  for  $d \in \{2, 3\}$ , and  $(1 + 2/\pi) \cdot (12k + 16)/\gamma$  for  $d \in \{4, 5\}$ .

## Experimental bound

One can see the result of [4] as a negative result, since we hoped that the Fréchet distance would be more robust under such projections. We performed the preliminary

---

<sup>1</sup>Note that the notion *upper bound* relates to the bound on  $1/c$ , used due to technical reasons.

experiments on the dataset of [6]. Their dataset  $\mathcal{D}$  contains 20199 realistic polygonal curves from  $\mathbb{R}^2$ , with complexities between 9 and 767. We have repeated the following procedure for 504 pairs of curves of  $\mathcal{D}$  selected uniformly at random. For each pair of curves (or their subcurves) the projection line was sampled  $r = 1000$  times. We observed the obtained distribution of the distortion  $c$  of the discrete Fréchet distance.

- i) We calculated the distortion  $c = d_F(P', Q') / d_F(P, Q)$  for the whole curves.
- ii) We observed the prefix curves  $P_\ell = \{p_1, \dots, p_\ell\}$  and  $Q_\ell = \{q_1, \dots, q_\ell\}$  of  $P$  and  $Q$  respectively, with complexity  $\ell$  equal 10, or to the multiples of 50. The distortion  $c = d_F(P'_\ell, Q'_\ell) / d_F(P_\ell, Q_\ell)$  is calculated.
- iii) For every prefix length  $\ell$  we chose at random subcurves of  $P$  and  $Q$  of complexity  $\ell$ , defined by  $\ell$  consecutive vertices of  $P$  and  $Q$  respectively. Let these curves be  $P_{\ell,r}$  and  $Q_{\ell,r}$ . We calculated the distortion  $c = d_F(P'_{\ell,r}, Q'_{\ell,r}) / d_F(P_{\ell,r}, Q_{\ell,r})$ .

This yielded 4286 pairs of (sub)curves. For these curves and their subcurves, the cumulative probability distributions of  $c$  were calculated, over the set of results of 1000 sampled runs. The Fréchet distance of the observed curves is not necessarily dominated by one pair of vertices, and varies upon which parts of the curves are observed. For all pairs of subcurves of  $P$  and  $Q$  and their respective projections  $P'$  and  $Q'$  we may assume that for any  $\gamma \in (0, 1)$  it is

$$\Pr[c \leq \gamma] \leq \gamma. \quad (2)$$

Indeed, when the cumulative probability distribution of the distortion  $c$  is observed over all tested pairs of curves, the mean and the standard deviation of the distortions obtained by our experiments for a given threshold  $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$ , suggest that for the realistic input curves  $P$  and  $Q$  the assumption of Equation (2) holds with high probability. The outlying maxima occur for the curves whose shape is similar to the curves from the proof of our lower bounds (see [4]), and thus strongly conditioned.

The cumulative probability distribution of the distortion  $c$  (Figure 1, left) suggests that for realistic input curves we can expect that  $\Pr[c \leq \gamma] \leq \gamma$ . This holds independently of the complexity  $t$  of the input curves, as illustrated by Figure 1 (right) for the given threshold  $\gamma = 0.5$ . This implies that with probability of at least 0.5 we expect that the discrete Fréchet distance will be reduced at most by a factor 2 when projected to a line chosen uniformly at random, independently of the input complexity. These results stand in stark contrast with our lower bounds. They indicate that highly distorted projections happen very rarely in practice, and only for strongly conditioned input curves.

Furthermore, it seems that the distortion of the discrete Fréchet distance is bounded by a constant (with high probability), and that it does not depend on the complexity  $t$  of the input curves, as suggested by Figure 1.

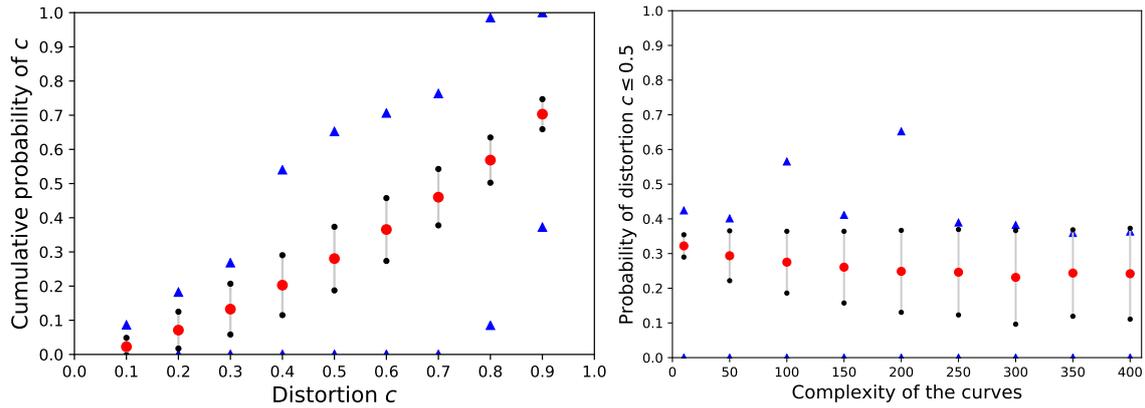


Figure 1: The cumulative probability distribution of the distortion (left). Given  $c \leq 0.5$ , the cumulative probability of distortion is shown as a function of the complexity  $t$  of the curves, for  $t \in \{10, 50, 100, 150, 200, 250, 300, 350, 400\}$  (right). The means  $\mu$  of the values denoted by red circles. The intervals  $[\mu - \sigma, \mu + \sigma]$  denoted by black dots, where  $\sigma$  is the standard deviation. The minima and maxima denoted by blue triangles.

## References

- [1] P. Afshani and A. Driemel. On the complexity of range searching among curves. In *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms*, pages 898–917, 2018.
- [2] K. Bringmann and M. Künnemann. Improved approximation for Fréchet distance on  $c$ -packed curves matching conditional lower bounds. *Int. J. Comput. Geom. Appl.*, 27(1-2):85–120, 2017.
- [3] A. Driemel, S. Har-Peled, and C. Wenk. Approximating the Fréchet distance for realistic curves in near-linear time. *Discrete & Computational Geometry*, 48(1):94–127, 2012.
- [4] A. Driemel and A. Krivošija. Probabilistic embeddings of the Fréchet distance. In L. Epstein and T. Erlebach, editors, *Proceedings of the 16th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 218–237, 2018.
- [5] A. Driemel, A. Krivošija, and C. Sohler. Clustering time series under the Fréchet distance. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 766–785, 2016.
- [6] M. Werner and D. Oliver. ACM SIGSPATIAL GIS Cup 2017: Range queries under Fréchet distance. *SIGSPATIAL Special*, 10(1):24–27, 2018. <http://sigspatial2017.sigspatial.org/giscup2017/home>.





Subproject A3  
Methods for Efficient Resource Utilization in  
Machine Learning Algorithms

Jian Jia Chen      Jörg Rahnenführer

# Feature Selection in High-Dimensional Data

Andrea Bommert  
Faculty of Statistics  
TU Dortmund University  
andrea.bommert@tu-dortmund.de

Feature selection is a key part of the analysis of high-dimensional data. In domains like bioinformatics, the models are not only used for prediction but they are also used for drawing biological conclusions which makes the interpretability and reliability of the models crucial.

For feature selection on high-dimensional datasets, filter methods play an important role, since they can be combined with any machine learning model and they can heavily reduce run time of machine learning algorithms. We have reviewed how different filter methods for feature selection work and compared their performance with respect to both run time and predictive accuracy. We conclude that there is no group of filter methods that always outperforms all other methods, but we make recommendations on filter methods, which perform well on many of the datasets. Also, we find groups of filters that are similar with respect to the order in which they rank the features.

In our current research, we work on selecting the correct subset of features for high-dimensional datasets with highly correlated features. We employ  $L_0$ -regularized regression and consider an adjusted stability measure during hyperparameter tuning.

## 1 Benchmark for Filter Methods for Feature Selection in High-Dimensional Data

Many methods for feature selection exist, but it is unclear which of these methods perform best. In our analyses we focused on the comparison of filter methods for feature selection.

We analyzed 22 filter methods that are available in the R package mlr [1] based on 16 high-dimensional datasets from various domains.

To find out which of the filter methods are similar with respect to the order in which they rank the features, we computed rank correlations, see Figure 1. We observed, that for

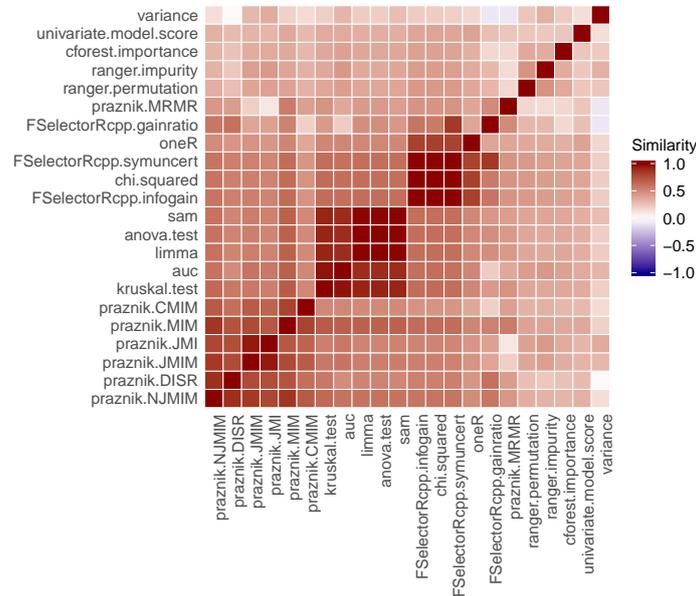


Figure 1: Rank correlations between the selection order of all pairs of filter methods on all datasets, averaged by the arithmetic mean. The filter methods are ordered by average linkage hierarchical clustering using the mean rank correlation as similarity measure.

some datasets, especially the datasets containing gene expression data with large number of features, there were three groups of similar filter methods and many filter methods which were not very similar to any other method. For the other datasets, most filter methods were very similar. The filters that were similar to each other mostly came from the same toolboxes. Also, we investigated the scaling behavior of the filter methods, identifying groups of filters which behave similarly with respect to run time.

Next, we analyzed the predictive accuracy of the features selected by the filter methods and the run time needed for feature selection and for building a good predictive model based on the selected features. We found out that there is no subset of filter methods which performs better than the rest of the filter methods on all datasets. Instead, the best filter methods, strongly differed between the datasets. Nevertheless, on average all filter methods performed better than not filtering at all. Also, the filter methods *ranger.permutation*, *ranger.impurity*, *FSelectorRcpp.symuncert*, *limma*, and *praznik.JMIM* performed well in many data situations and therefore seem advisable, see Figure 2. The filters *FSelectorRcpp.symuncert*, *limma*, and *praznik.JMIM* can be seen as representatives of the three groups of similar filter methods. Filters

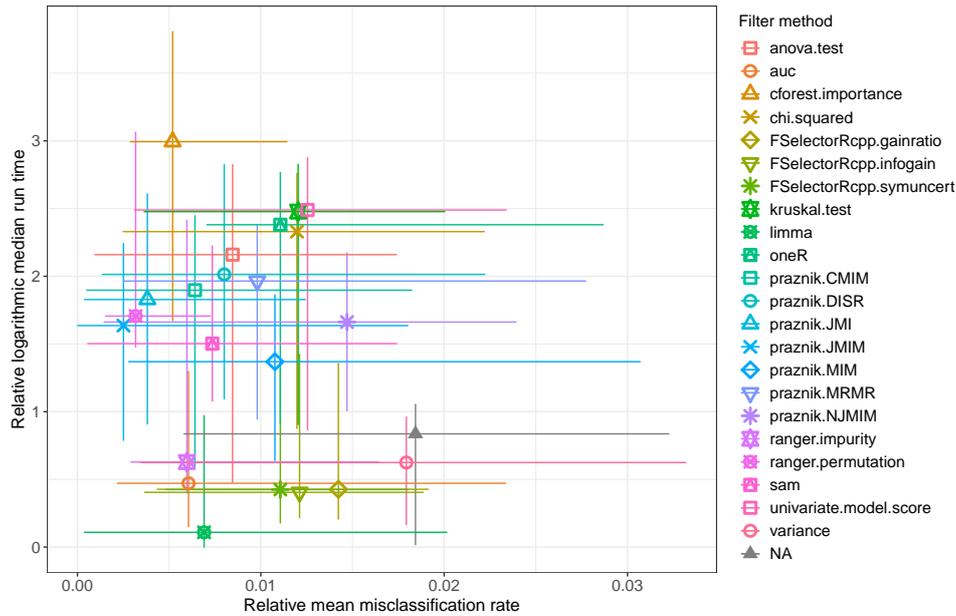


Figure 2: Relative mean misclassification rate and relative logarithmic median run time of the filter methods with optimal configurations aggregated over all datasets. The median of both performance measures (relative mean misclassification rate and relative logarithmic median run time) across all datasets is displayed by a symbol. The upper and lower quartile are located at the respective ends of the horizontal and vertical lines.

*ranger.impurity*, *FSelectorRcpp.symuncert*, and *limma* achieved very low run times. For filter *ranger.permutation* we observed comparably high run times but also very high predictive accuracy. So, if only limited computational resources are available for finding a suitable filter method, we recommend trying *ranger.permutation*, *ranger.impurity*, *FSelectorRcpp.symuncert*, *limma*, and *praznik.JMIM*. However, whenever possible one should try to find the best filter method for a given dataset. This can be done by considering it as a tuning parameter. This work is submitted but not yet published.

## 2 Selecting the Correct Subset of Features by Considering an Adjusted Stability Measure

In our current research, we take into account the fact that for high-dimensional datasets (e.g. microarray data) there often are highly correlated features. In [2], we proposed performing the hyperparameter tuning of a machine learning method (which includes feature selection in the model fitting process) with respect to both predictive accuracy and feature selection stability. For datasets with many highly correlated features, the stability

measures which are commonly used seem unsuitable: They rate it as unstable if instead of one feature, another feature - which is almost identical to the first feature but has a different name - is chosen. For the context of gene lists, adjusted stability measures exist which take into account feature similarities [4]. We generalize one of these measures so that it is applicable in the context of feature selection.

Because we want to perform feature selection on datasets with highly correlated features, we should use a feature selection method which does not select all highly correlated features of a relevant feature. We choose to employ  $L_0$ -regularized logistic regression which possesses this property unlike most of the other classification methods with embedded feature selection. There is one hyperparameter that needs to be tuned.

We compare the approaches (1) tuning w.r.t. predictive accuracy, (2) tuning w.r.t. predictive accuracy and an unadjusted stability measure, (3) tuning w.r.t. predictive accuracy and an adjusted stability measure, and (4) stability selection [3] with parameters tuned w.r.t. predictive accuracy, based on both simulated and real data. On simulated data, we conclude, that using an adjusted stability measure helps with selecting the correct features (the features that have been used for generating the values of the target variable). In contrast to the approaches (1) and (2), fewer redundant features are selected. Compared to stability selection, fewer relevant features are forgotten. On real datasets, we can only evaluate the predictive performance and the sparsity of the models. We observe that the models fitted with our proposed approach are of comparable predictive quality as the models of the established approaches and for many datasets, they are more sparse. This work is still in progress but the preliminary results are promising.

## References

- [1] Bernd Bischl, Michel Lang, Lars Kotthoff, Julia Schiffner, Jakob Richter, Erich Studerus, Giuseppe Casalicchio, and Zachary M. Jones. mlr: Machine learning in R. *Journal of Machine Learning Research*, 17(170):1–5, 2016.
- [2] Andrea Bommert, Jörg Rahnenführer, and Michel Lang. A multicriteria approach to find predictive and sparse models with stable feature selection for high-dimensional data. *Computational and Mathematical Methods in Medicine*, 2017:Article ID 7907163, 2017.
- [3] Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- [4] Min Zhang, Lin Zhang, Jinfeng Zou, Chen Yao, Hui Xiao, Qing Liu, Jing Wang, Dong Wang, Chenguang Wang, and Zheng Guo. Evaluating reproducibility of differential expression discoveries in microarray studies by considering correlated molecular changes. *Bioinformatics*, 25(13):1662–1668, 2009.

# Model-Based Optimization on Parallel Infrastructures and Functions with Concept Drift

Jakob Richter  
Faculty of Statistics  
TU Dortmund University  
richter@statistik.tu-dortmund.de

Optimization of parametrized algorithms with a long runtime is a challenging task. The goal is to find a parameter configuration that leads to an optimal performance of the algorithm, whereas the performance is measured by an arbitrary criterion. Therefore we need to evaluate the algorithm on a set of parameter configurations. As the runtime of each algorithm evaluation is very long the total number of algorithm evaluations that can be calculated within an acceptable time bound is limited. The optimization algorithm has to use the limited budget as efficiently as possible. Model-based optimization (MBO) is a popular technique for such expensive black-box optimization. By fitting a regression model on the set of already evaluated parameter configurations it is possible to obtain a prediction of the performance for new parameter configurations. This so called surrogate is used to guide the optimization to promising areas in an iterative fashion. An infill criterion determines the next parameter configurations to be evaluated on the expensive black-box. The result is then used to improve the surrogate. Using this technique we aim to solve two problems: 1. Evaluate multiple parameter configurations in parallel if the infrastructure is available and handle heterogeneous runtimes to avoid idling using our RAMBO-Framework. 2. Apply MBO to optimize a black-box over time and handle concept drifts, where the relation of the parameter configurations and the performance changes over time. Applications reach from hyperparameter tuning for machine learning methods to automatic pipeline optimization.

# 1 MBO on Parallel Infrastructures

Ordinary MBO is sequential by design. In each iteration a single parameter configuration is proposed and evaluated. The need to evaluate more configurations within the same time and the growth of parallel computer infrastructure have driven the development of MBO methods that propose multiple parameter configurations [4, 5] in each iteration to be evaluated in parallel. However, they neglect the fact that the runtime of the black-box often depends on the parameter configuration. If the parallel MBO method is configured to propose  $k$  parameter configurations for  $k$  available workers in each iteration, then each worker evaluates a single configuration. This can lead to idling if evaluation times differ. In those cases, the workers will wait after they completed a evaluation until new proposals are generated. However, these proposals are only generated after all evaluations are finished. Idling is a problem, because it means that time is wasted on idling workers. This time could be used to evaluate other parameter configurations. This strictly iterative procedure is called *synchronous parallelization*.

We tackle this problem by introducing another regression model that predicts the runtime of a parameter configuration based on the runtime measured on previous algorithm evaluations, similarly to the performance prediction of the surrogate model. This allows us to obtain runtime predictions for the proposed parameter configurations, and consequently calculate expected idling. Our proposed method prevents idle times by proposing more than  $k$  parameter configurations and using knapsack scheduling to select the best subset. This method is described in detail in Kotthaus et al. [6]. There, we also compare our proposed method to *asynchronous* approaches. In the asynchronous approach, each worker runs independently, and a new proposal is generated on the bases of all completed evaluations and the pending evaluations [3] as soon as the evaluation is finished. Hence idling does not occur anymore. However the computational overhead is increased, because outcomes of pending evaluations have to be estimated and each worker has to construct its own surrogate.

# 2 MBO on Functions with Concept Drift

Typically for Model-based optimization, the objective function is assumed to be constant over time. However, time-varying functions are found in many domains and issue a new challenge for the MBO framework. We developed two novel adaptations of the model-based optimization framework to handle such concept drifts. The first adaptation uses a time window on the evaluations that the surrogate is trained on; the second includes the time directly as an additional covariate in the surrogate model. We compare our adaptations to ordinary MBO on scenarios with no, gradual and sudden concept drifts, as shown in Figure 1.

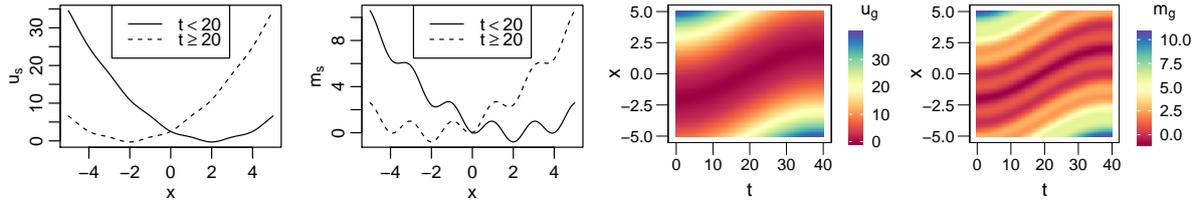


Figure 1: *Left:* Functions  $u_s$  and  $m_s$  with sudden concept drift at  $t = 20$ . *Right:* Functions  $u_g$  and  $m_g$  with gradual concept drift.

First results (Figure 2) show that in presence of a concept drift, the novel approaches using a time window ( $\tau w5$ ,  $\tau w10$ ) and including the time as covariate in the surrogate ( $\tau ac$ ) yield clearly better results than the uncorrected optimizers. If no concept drift is present, the adaptations seem to only decrease the performance slightly.

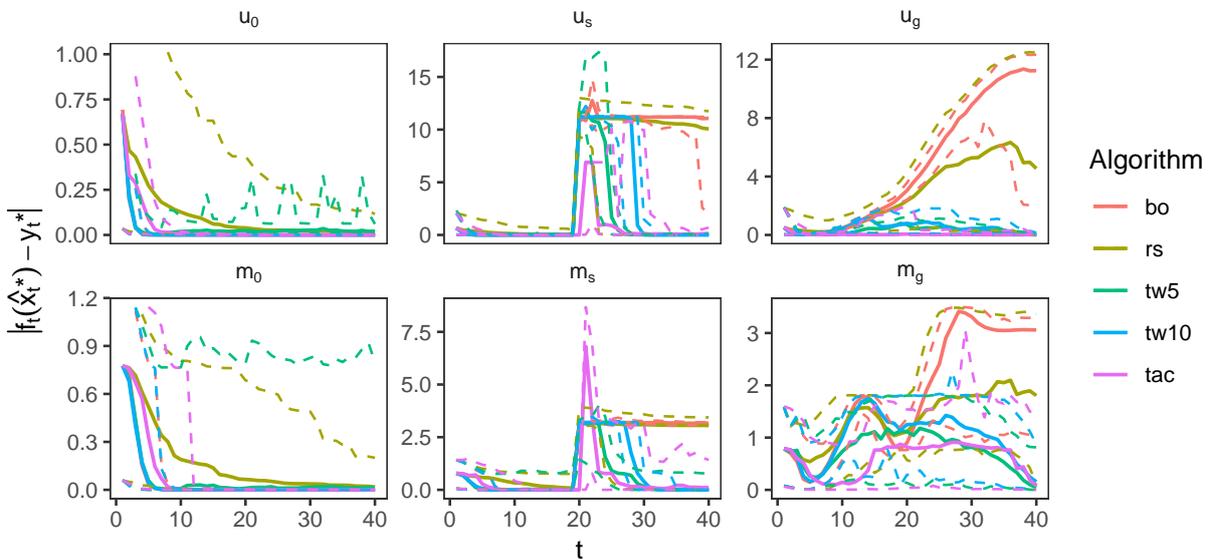


Figure 2: Difference of the theoretical optimal performance at time  $t$  and the performance obtained by the proposed optimum of the optimization strategy. *Top:* Unimodal functions  $u_0$  with no,  $u_s$  with sudden and  $u_g$  with gradual concept drift. *Below:* Multi-modal functions with same drifts. A smaller window ( $\tau w5$ ) leads to a faster adaption ( $u_s$ ,  $m_s$ ) but also increases uncertainty ( $u_0$ ,  $m_0$ ).

### 3 Further Research

I participate actively in the development of the R-Package **mlr** [1], it's successor **mlr3**<sup>1</sup> and **mlrMBO** [2] amongst others. The latter is the foundation of the RAMBO Frame-

<sup>1</sup><https://github.com/mlr-org/mlr3>

work and the MBO-CD algorithms. We developed the *OpenML TuneBench Server*<sup>2</sup>. It uses a database of a vast amount of pre-evaluated hyperparameter configurations for a combination of machine learning methods and datasets [7]. This allows us to create a benchmark suite for parameter optimization that does not require the actual costly evaluation of the machine-learning algorithms. As a result we are able to benchmark out methods more efficiently on real world problems.

## References

- [1] Bernd Bischl, Michel Lang, Lars Kotthoff, Julia Schiffner, Jakob Richter, Erich Studerus, Giuseppe Casalicchio, and Zachary M. Jones. “Mlr: Machine Learning in R”. In: *Journal of Machine Learning Research* 17.170 (2016), pp. 1–5.
- [2] Bernd Bischl, Jakob Richter, Jakob Bossek, Daniel Horn, Janek Thomas, and Michel Lang. “mlrMBO: A Modular Framework for Model-Based Optimization of Expensive Black-Box Functions”. In: *arXiv:1703.03373 [stat]* (Mar. 2017), pp. 1–26.
- [3] David Ginsbourger, Janis Janusevskis, and Rodolphe Le Riche. *Dealing with Asynchronicity in Parallel Gaussian Process Based Global Optimization*. Tech. rep. 2011, pp. 1–27.
- [4] David Ginsbourger, Rodolphe Le Riche, and Laurent Carraro. “Kriging Is Well-Suited to Parallelize Optimization”. In: *Computational Intelligence in Expensive Optimization Problems*. Springer, 2010, pp. 131–162.
- [5] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. “Parallel Algorithm Configuration”. en. In: *Learning and Intelligent Optimization*. Ed. by Youssef Hamadi and Marc Schoenauer. Lecture Notes in Computer Science 7219. Springer Berlin Heidelberg, 2012, pp. 55–70. DOI: 10.1007/978-3-642-34413-8\_5.
- [6] Helena Kotthaus, Jakob Richter, Andreas Lang, Janek Thomas, Bernd Bischl, Peter Marwedel, Jörg Rahnenführer, and Michel Lang. “RAMBO: Resource-Aware Model-Based Optimization with Scheduling for Heterogeneous Runtimes and a Comparison with Asynchronous Model-Based Optimization”. In: *Learning and Intelligent Optimization*. Lecture Notes in Computer Science. Springer, Cham, June 2017, pp. 180–195. DOI: 10.1007/978-3-319-69404-7\_13.
- [7] Daniel Kühn, Philipp Probst, Janek Thomas, and Bernd Bischl. “Automatic Exploration of Machine Learning Experiments on OpenML”. In: *arXiv:1806.10961 [cs, stat]* (June 2018), pp. 1–6.

---

<sup>2</sup><https://github.com/mboecker/lookup-server>





Subproject A4  
Resource efficient and distributed platforms  
for integrative data analysis

Michael ten Hompel

Christian Wietfeld

# LoRaWAN Networks and its capabilities to contribute to 5G massive Machine Type Communications

Stefan Böcker

Lehrstuhl für Kommunikationsnetze

Technische Universität Dortmund

stefan.boecker@tu-dortmund.de

Internet of Things applications, interconnecting our daily life, have received a certain amount of interest. A huge number of potential technology solutions are available, but a comprehensive networking solution based on one technology seems unlikely. The development of 5G mobile networks with a desired node density of 1,000,000 devices per square kilometer in the area of *massive Machine Type Communication* (mMTC) faces a growing number of unlicensed technologies enabling a simple, cost-effective network operation independent of licensed operators. Specifically, this work aims to analyze the suitability of LoRaWAN to contribute to given 5G requirements for specific mMTC applications in large-scale deployments. The performance evaluation illustrates that LoRaWAN is attractive due to high communication ranges up to multiple kilometers, enabling a high coverage even with a small number of cells. However, due to a simple channel access mechanism in combination with regulatory requirements defined for the 868 MHz short-range device (SRD) frequency band, this is accompanied by low data rates and high delays of several seconds in large-scale scenarios. Nevertheless, LoRaWAN indicates a high potential to contribute to 5G mMTC application areas, especially for non-time-critical sensor applications.

## 1 Introduction

Internet of Things (IoT) applications have experienced significant attention and growth related to large scale deployments, establishing smart object and sensor networks, covering use cases of industrial, commercial, public as well as private domains. The idea of

new basic technologies for interconnecting urban infrastructure, building management, private and industrial transportation, as well as energy supply is not only discussed in research, but also increasingly implemented in real deployments. In particular, the fifth mobile radio generation (5G) promises stable connectivity for a very high number of devices located in small areas. The ITU-R defines related requirements for several Internet of Things applications within the area of (*massive Machine Type Communication* (mMTC) [1]). In particular technologies in licensed frequency bands, tailored to these specific IoT requirements, e.g. LTE evolutions Narrowband-IoT (NB-IoT) or eMTC, but also counterparts within unlicensed frequency bands, so-called Low Power Wide Area Networks (LPWAN), are discussed as major technology solutions. In order to stress especially technologies from the unlicensed sector, this work focuses on scalability analysis of LoRaWAN networks based on typical application activities and interference scenarios.

## 2 Approach and Implementation

As illustrated in Figure 1 the LoRaWAN performance evaluation is addressing 5G mMTC requirement definitions, while at the same time mandatory regulatory restrictions to guarantee a fair, harmonized operation in unlicensed frequency bands, are taken into account.

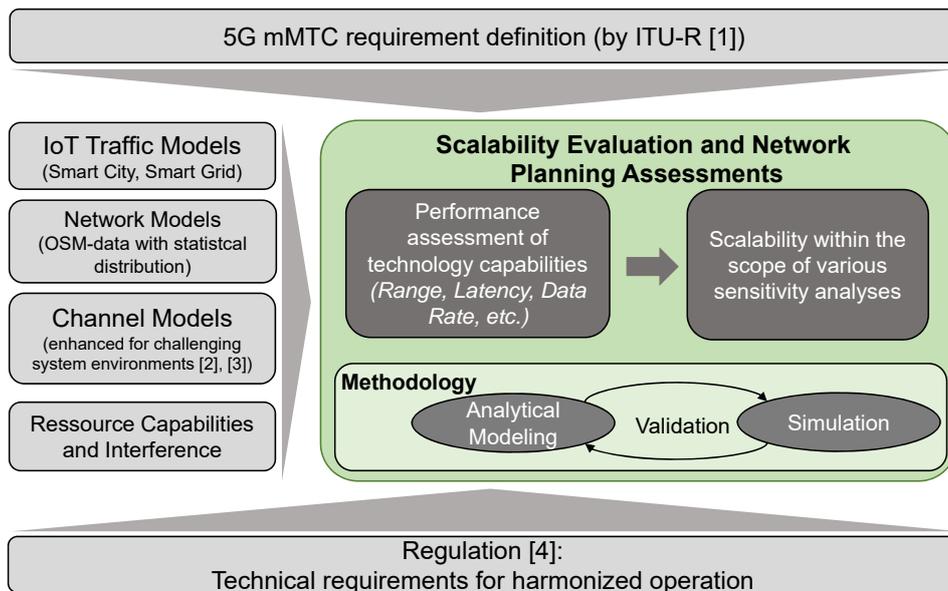


Figure 1: Modeling approaches for LoRaWAN scalability analysis

The performance evaluation itself is performed by two independent modeling approaches. First, an analytical model based on current state-of-the-art [6] is evolved to analyze the key performance indicators such as range, latency, data rate etc. Major enhancements are detailed latency evaluations even for large-scale scenarios, which improve scalability results by means of prediction of service qualities. In order to verify the analytical model, this work utilizes a simulation model [5] covering LoRaWAN channel access, which is enhanced by necessary functionalities to cover former introduced latency and throughput results for large-scale deployments. The reliability of both models is verified within a cross-validation process, whereby both models enable sensitivity analysis for traffic amount, duty cycle, as well as resource utilization. Necessary input is detailed in specific traffic models focused on IoT applications: network models covering real world

data and statistical distributions, a channel model fitted to urban LoRaWAN deployments and challenging system environments, as well as a model that provides the opportunity to consider specific resource and interference situations.

### 3 Performance and Scalability Evaluation

Our performance evaluation is performed for several activity levels, variations of considered payload size and channel configurations. Results for an exemplary parameter constellation (maximum duty cycle of 1% with a considered payload size of 32 Byte) are presented in Figure 2. The maximum throughput (uplink) is illustrated for each mandatory, orthogonal data rate class (DR0-5), resulting in the aggregated system throughput, which is depicted with and without a capture effect.

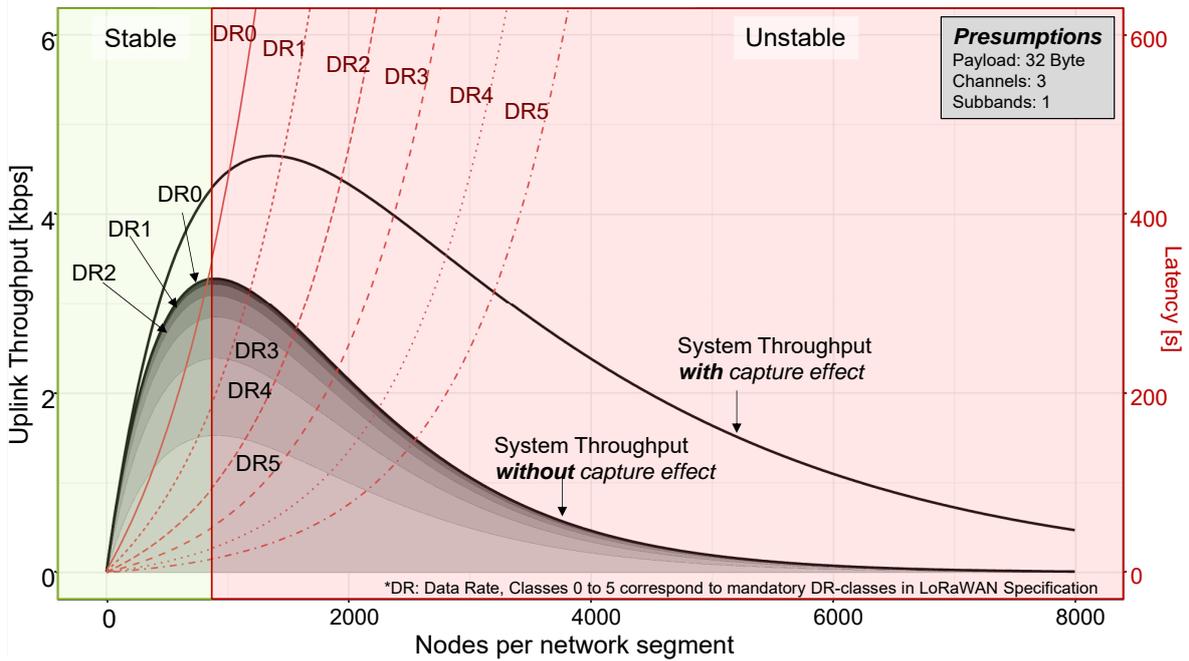


Figure 2: LoRaWAN Throughput utilizing maximum Duty Cycle

It can be seen that a maximum system throughput (uplink) of 3.25 kbps is achieved for about 950 nodes per network segment. In addition, correlating latencies for each data rate class are illustrated. In order to derive a certain service quality, it is assumed that the maximum latency is always given by the lowest data rate class (DR0). Thus, the expected average latency can be determined with up to 320 seconds. Considering the capture effect, the maximum system throughput is increased to 4.5 kbps generated by about 1500 nodes, while at the same time, the correlated latency remains stable. Considering a fairly lower inter-arrival time (IAT) of 2 hours compared to Figure 2 and an increased system capacity (5 channels each 125 kHz) the maximum scalability of LoRaWAN networks in the 868 MHz can be increased to about 175.000 nodes per network segment.

### 4 Conclusion and Further Research

In this work, a performance analysis approach in order to evaluate the suitability of LoRaWAN networks to contribute to the given 5G mMTC requirement of about 1.000.000

devices per square kilometer is presented. Scalability results verify that LoRaWAN is a feasible technology solution that can contribute to the future 5G IoT area, whereby the application field should be limited to non-time-critical sensor applications. Currently, the LoRaWAN evaluation is limited to analytical methods cross-verified by means of simulations. This can be enhanced by lab and field trials, in order to further validate and strengthen already achieved evaluation results. In future work, the LoRaWAN evaluation will be extended to LoRaWAN implementations at 2.4 GHz, which supports significantly larger bandwidths and thus promises increased scalability.

## References

- [1] ITU, "Minimum requirements related to technical performance for IMT-2020 radio interface(s)," International Telecommunication Union, Tech. Rep. ITU-R M.2410-0, 2017.
- [2] P. Jörke, S. Böcker, F. Liedmann, and C. Wietfeld, "Urban channel models for smart city IoT-networks based on empirical measurements of LoRa-lmks at 433 and 868 MHz," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, PIMRC*, vol. 2017-Octob, 2018.
- [3] S. Monhof, S. Böcker, J. Tiemann, and C. Wietfeld, "Cellular Network Coverage Analysis and Optimization in Challenging Smart Grid Environments," *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 2018.
- [4] ETSI, *Short Range Devices (SRD) operating in the frequency range 25 MHz to 1 000 MHz; Part 2: Harmonised Standard covering the essential requirements of article 3.2 of Directive 2014/53/EU for non specific radio equipment*, European Telecommunications Standards Institute Std. EN 300 220-2, 2016.
- [5] M. Slabicki, G. Premsankar, and M. D. Francesco, "Adaptive configuration of lora networks for dense IoT deployments," in *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*, apr 2018, pp. 1–9.
- [6] K. Mikhaylov, J. Petäjajarvi, and T. Hänninen, "Analysis of Capacity and Scalability of the LoRa Low Power Wide Area Network Technology," *European Wireless 2016*, pp. 119–124, 2016.

# Workload Adaptation for Energy Harvesting Driven, Battery-Free Systems

Markus Buschhoff

Department of Computer Science 12

Technische Universität Dortmund

markus.buschhoff@tu-dortmund.de

The vision of Industry 4.0 demands for distributed, autonomous systems to drive logistics and production facilities. A large set of these systems is hard to maintain due to their energy demands, charging cycles and battery decay. To suit the needs for battery-free Industry 4.0 appliances, inexpensive energy harvesting approaches are required that adapt a node's performance and energy consumption to its energy income. In an experimental setup, a new method was evaluated to show an inexpensive approach in lab conditions.

## 1 Introduction

In-situ measurement for mobile devices often requires a voltage supply that is reliable to a certain degree. This is necessary to generate a stable reference voltage for the measurement and to drive components like analog-digital converter or measurement amplifiers. Although classical approaches yield good accuracies, both the energetic requirements and the expenses of components may not be acceptable for harvesting-driven Industry-4.0 applications.

The approach used here takes advantage of simple components. To achieve this, a software-driven solution was created that measures the transition time of a general purpose IO pin of a microcontroller. It requires and abuses some internal circuitry of the microcontroller (pull resistors and gate capacity) and a few, inexpensive external components, namely a z-diode and a resistor.

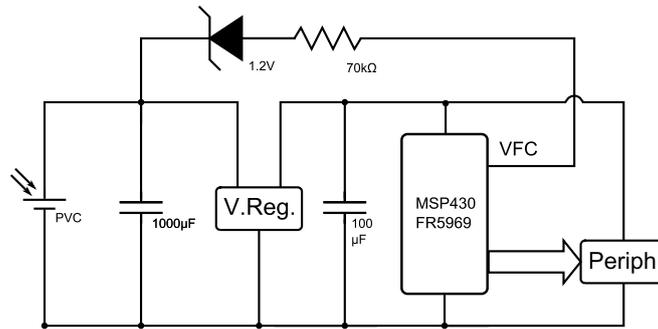


Figure 1: Basic setup of a node with photovoltaic cell (PVC) and in-situ measurement deployment.

## 2 System Setup

Our example system, as shown in Figure 1, consists of a photovoltaic cell (PVC) to harvest energy, and a capacitor that is able collect enough charge for a given hardware task. The capacitor should be dimensioned with reserves (factor 2) to compensate for inaccuracies and energy loss. The capacitor feeds a voltage regulator (VREG) to hold a steady supply voltage for the microcontroller. The capacitor charge is estimated by the following measures:

1. An input pin of the microcontroller is reset to 0 state by activating the internal pull down resistor in software.
2. The capacitor voltage is reduced to a value close above the threshold voltage of the IO pin by using an Z-diode.
3. The current is reduced to achieve a minimal charge transfer by using a resistor in the high  $k\Omega$  range.
4. The time to achieve a low/high transition on the IO pin is measured by the microcontroller. The charge state of the capacitor is derived from this value.
5. Depending on the charge state, a peripheral hardware workload is triggered. Independent of this, the system idles for 10 ms.

In practice, step 4 can be easily achieved in hardware by using a capture input of a microcontroller that is able to trigger a timer capture and/or an interrupt on a high flank.

In our setup, a timer value was captured every 10 ms, and the system decides to either start a workload at this time or to idle. The decision is impacted by preceding decisions, so that in integrating regulator loop is achieved: When low on voltage (decided by a threshold value), a pause variable  $p$  is increased by a fraction  $R$  of the measured capture time  $c$ . Otherwise,  $p$  will be decreased. After a cycle of subsequent sufficient voltage

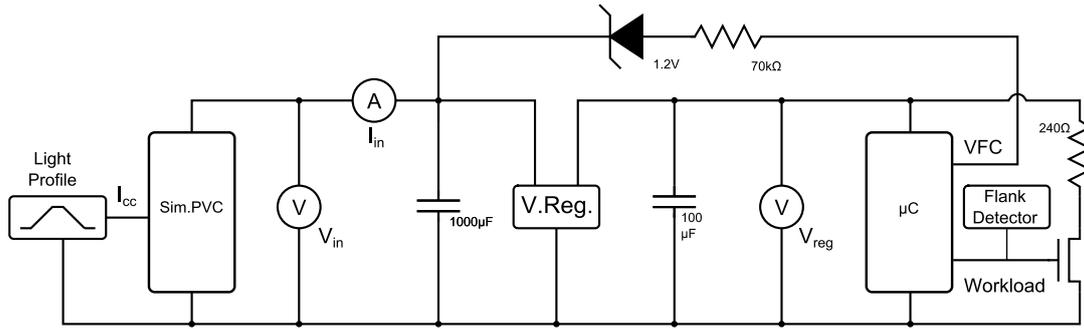


Figure 2: Experiment setup

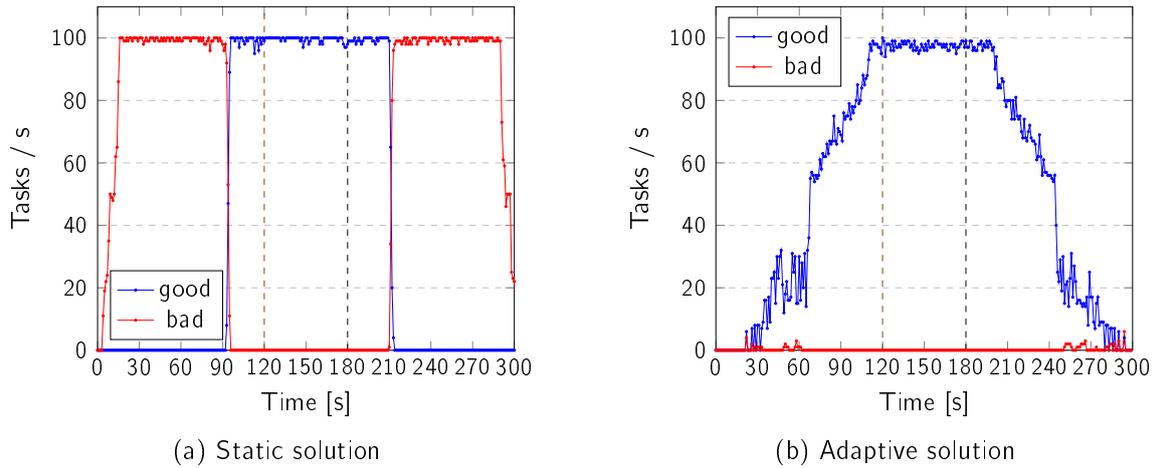


Figure 3: Good and bad tasks: (a) no adaptation (b) proposed solution

conditions,  $p$  will eventually decrease to zero. It can now be assumed that the capacitor was charged long enough to achieve a complete hardware task.

### 3 Experiment

To establish an evaluation of this approach, a photovoltaic cell (PVC) emulation was used to generate a well defined energy income. The whole setup is depicted in Figure 2. The emulated PVC runs from zero output to maximum output linearly over 2 minutes, resides for one minute at peak power production and then decreases to zero within the next 2 minutes. During this time, the consumed energy was measured at the PVC and at the peripheral hardware workload. The task of the system is to drive a periphery device which is simulated by a resistor being switched on and off by a transistor for 10 ms. The supply voltage of the workload was continuously observed. If it stays within a boundary of 10% of the nominal supply voltage (3 V), the task is considered as complete (marked as good), otherwise as interrupted by a power-fail (marked as bad). Figure 3 shows

the results for two experiment runs. The left side shows a run with no measurement deployment and a static schedule of starting the task at each available time slot. The right side shows the proposed solution, showing a surplus of good tasks during the dusk and dawn period of the daylight simulation, while the number of bad stays close to zero during the whole experiment.

	good	bad	effectivity	efficiency
adaptive	16375	62	0.693	0.905
static-100	11669	16635	0.493	0.879
static-50	10056	4407	0.425	0.644

Table 1: Task counters, efficiency and accuracy.

As a metric for this approach, *energy efficiency* is calculated by the ratio of consumed energy at the PVC and the (known) total available energy at the PVC-emulator's output. Efficiency here is a measure for good utilization of a PVC, which has to be loaded at electrical conditions near its *maximum power point* (MPP) to be efficient. The MPP has to be tracked continuously, as it changes with light conditions. Usually, this is done by external hardware, so called MPP-Trackers or MPPTs, which is not necessary here.

Efficiency does not reflect the purpose for which energy was used, e.g. it could have been wasted completely for measurement, while not supplying tasks with sufficient voltage. Thus, a metric of *effectivity* is introduced, which is the ratio between the actual amount of good tasks, and the theoretically possible good tasks, when no energy loss for the microcontroller or energy measurement occurs.

Table 1 shows some of the results for three experiments. The adaptive solution as proposed, a solution called *static-100*, which uses every time slot to start a task, and a solution called *static-50* using every second slot and idling during the rest of the time.

## 4 Discussion

The proposed approach has proven to be well able to adapt its energetic workload to environmental harvesting conditions. It significantly reduces the number of peripheral hardware tasks that are interrupted due to low power conditions, while at the same time raising the number of good tasks by 40%. A static solution can be created to mimic these results in a single scenario, but cannot adapt to changes in harvesting conditions.

The adaptive solution also has shown to be capable of satisfying the maximum power point conditions of a PVC harvester and shows an efficiency, i.e. tracking accuracy, of 90%. It renders the implementation of expensive hardware MPPT circuits useless for this kind of scenario. Finally, 70% of the available energy was used to run good tasks. This is in the same order as other, more complex and expensive solutions in the field.

# Machine Learning Based Transmission Power Prediction in LTE User Equipment

Robert Falkenberg  
Lehrstuhl für Kommunikationsnetze  
Technische Universität Dortmund  
robert.falkenberg@tu-dortmund.de

The design of energy-aware devices for the Internet of Things (IoT), such as sensor nodes in vehicles or mobile robotic systems, requires knowledge about the power consumption of the communication system. In cellular networks like Long Term Evolution (LTE) or 5G uplink, the lion's share of the overall power consumption is caused by uplink transmissions and highly depends on the required transmission power. However, on most of today's commercial User Equipment (UE), this information is not accessible at the application layer. Moreover, even most of the established simulation frameworks do not explicitly model the transmission power control mechanisms. In this report, we summarize our findings in closing those gaps by applying machine learning on empirical data sets from drive tests in a public LTE network and provide models which predict the transmission power of a mobile device.

## 1 Introduction

Due to the steadily increasing number of energy-constraint cellular IoT devices, energy-aware design of systems, networks and applications becomes a more and more important challenge as it directly affects the operation time in the field. A large fraction of the available power resources are spent by the communication system, especially by uplink transmissions and the involved transmission power [2], [4].

Although the exact value of the transmission power is well known to the device in the instant of an ongoing transfer, it is not exposed to the application layer (e.g. by an Application Programming Interface (API)) on most mobile handsets, as shown in Fig. 1. This circumvents cross-layer energy-aware applications, which avoid high-power

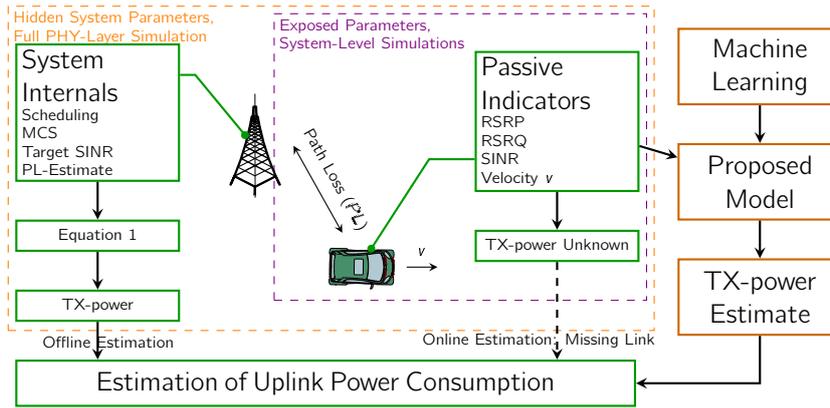


Figure 1: Overview of hidden and available network indicators in LTE User Equipment (UE) and the relationship to the transmission power and power consumption.

transmissions, e.g., by delaying or redirecting to a different link. Furthermore, channel-aware transmission schemes, such as ML-CAT [7], can greatly benefit of knowledge about the transmission power requirements in terms of an additional input value or as an output value to evaluate the impact of the transmission scheme on the power consumption [8]. In addition, even most established LTE simulators like SimuLTE [9] and LTE-Sim [6] do not consider precise power consumption models but stick to very simple linear approximations at system-level.

However, both, mobile handsets and simulators provide at least a minimum of information about the strength and quality of the downlink signal, which in turn also affects the power control mechanism in uplink direction (cf. next section). Knowledge about velocity, amount of exchanged data and the actual data rate provide further conclusions about the involved transmission power.

In this report, we present a machine learning-based approach to predict or estimate the transmission power of LTE UE in system-level simulations and on Commercial Off-the-Shelf (COTS) smartphones based on the limited set of available indicators in these scenarios. The interested reader is kindly referred to [3] for a deeper insight into this topic.

## 2 Approach and Setup

According to the LTE standard [1], the UE transmission power  $P_{tx}$  corresponds

$$P_{tx} = \min(P_{max}, P_0 + 10 \log_{10}(M) + \alpha \cdot PL + \Delta_{MCS} + \delta). \quad (1)$$

This includes the estimated path loss  $PL$ , a pre-configured Fractional Path Loss Compensation (FPC) factor  $\alpha$ , and the requested minimum Signal to Interference and Noise Ratio (SINR)  $P_0$  for signal reception at the base station. While  $PL$  correlates with the most commonly available indicator Reference Signal Received Power (RSRP),  $\alpha$  and  $P_0$  are operator-specific constants for specific environments, e.g. for rural or urban regions. The number of allocated Resource Blocks (RBs)  $M$  and the Modulation and Coding

Table 1: Captured Features and Association to Application-Specific Prediction Models based on Full-Feature Set  $\mathbb{F}$ , Practical Sets  $\mathbb{P}1/\mathbb{P}2$ , and Simulation Set  $\mathbb{S}$ .

Parameter	Model	Indicated Influences(s)
Velocity	$\mathbb{F}, \mathbb{P}1, \mathbb{P}2, \mathbb{S}$	Distortions by fast fading
Upload size	$\mathbb{F}, \mathbb{P}1, \mathbb{P}2, \mathbb{S}$	Influence of TCP slow start
RSRP	$\mathbb{F}, \mathbb{P}1, \mathbb{P}2, \mathbb{S}$	Signal strength, distance
RSRQ, SINR	$\mathbb{F}, \mathbb{P}1, \mathbb{P}2$	Signal clarity, interference
Datarate	$\mathbb{F}, \mathbb{P}1$	Signal strength, allocated RBs $M$
RSSI	$\mathbb{F}$	Signal strength, distance
Frequency band	$\mathbb{F}$	Environment [5]
Number of neighbor cells (intra/inter freq.)	$\mathbb{F}$	Environment, cell density, interference
Cell bandwidth	$\mathbb{F}$	Exhaustion of TX-power headroom

Scheme (MCS) specific offset  $\Delta_{MCS}$  directly influence the uplink data rate of the device. The closed-loop component  $\delta$  reflects feedback information to fulfill the signal quality requirements of the base station and should average to 0 for a well-configured cell. Finally, the power is capped by the maximum transmission power  $P_{max}$  of the device.

On this foundation, we performed field measurements in the LTE network of an operator in Germany. Along a trajectory of 44 km between the two cities Dortmund and Hamm, we instructed an embedded Vehicle-to-everything (V2X) platform to perform periodic data uploads and to log 31 parameters (e.g. RSRP, Reference Signal Received Quality (RSRQ), velocity, data rate, upload size) in intervals of 1 s [3]. In contrast to COTS smartphones, the V2X platform also reports the instant transmission power of the modem, which is used as a label in the subsequent machine learning process.

### 3 Machine Learning Results

Based on the 6172 samples collected during the campaign, we applied three methods of machine learning to train prediction models.

This comprises Ridge Regression, Random Forest, and Deep Learning, each including 10-fold cross-validation. Besides a reference model on the full feature set ( $\mathbb{F}$ ) on all parameters in Tab. 1, we also trained the models on feature subsets which fit practical applications on mobile handsets ( $\mathbb{P}1$  and  $\mathbb{P}2$ ) and system-level simulations ( $\mathbb{S}$ ).

The accuracy of the trained models in terms of Root Mean

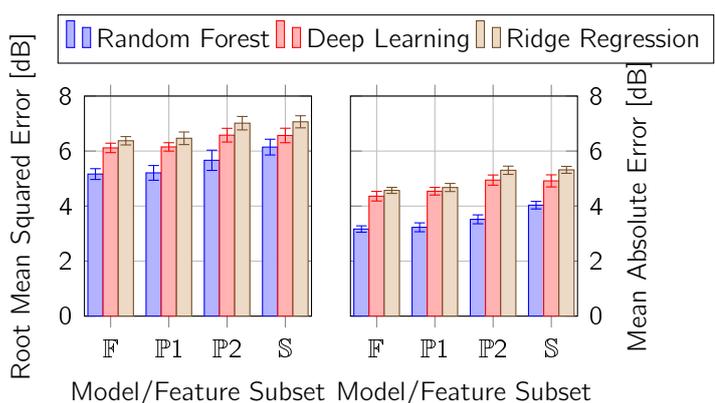


Figure 2: Prediction error of the transmission power.

Square Error (RMSE) and Mean Absolute Error (MAE) are shown in Fig. 2. Error bars are displaying the standard deviation computed over the cross-validation runs, which are well below 0.4 dB and indicate a good model fit to unknown data. The best results are achieved with Random Forest in the full-featured model  $\mathbb{F}$  with an MAE of 3.166 dB. However, even in the smallest feature subset  $\mathbb{S}$ , the error raises only moderately by less than 1 dB to 4.033 dB. The full analysis of this approach is documented in [3].

## 4 Conclusion and Further Research

With the power of machine learning, we presented a method to predict and estimate the transmission power of LTE UE with an error in the range of 3 dB to 4 dB. The computation is based on the little subset of indicators available on COTS devices and system-level simulators, thus enables in the next step an online application of the Context-Aware Power Consumption Model (CoPoMo) [2].

## References

- [1] 3gpp ts 36.213 - physical layer procedures (release 13), January 2016.
- [2] Björn Dusza, Christoph Ide, Liang Cheng, and Christian Wietfeld. CoPoMo: A Context-Aware Power Consumption Model for LTE User Equipment. *Transactions on Emerging Telecommunications Technologies (ETT)*, Wiley, 24(6):615–632, October 2013.
- [3] Robert Falkenberg, Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. Machine learning based uplink transmission power prediction for lte and upcoming 5g networks using passive downlink indicators. In *2018 IEEE 88th IEEE Vehicular Technology Conference (VTC-Fall)*, Chicago, USA, Aug 2018.
- [4] Robert Falkenberg, Benjamin Sliwa, and Christian Wietfeld. Rushing full speed with lte-advanced is economical - a power consumption analysis. In *IEEE Vehicular Technology Conference (VTC-Spring)*, Sydney, Australia, jun 2017.
- [5] P. Joshi, D. Colombi, B. Thors, L. E. Larsson, and C. Törnevik. Output power levels of 4G user equipment and implications on realistic RF EMF exposure assessments. *IEEE Access*, 5:4545–4550, 2017.
- [6] G. Piro, L. A. Grieco, G. Boggia, F. Capozzi, and P. Camarda. Simulating LTE cellular systems: An open-source framework. *IEEE Transactions on Vehicular Technology*, 60(2):498–513, 2011.
- [7] Benjamin Sliwa, Thomas Liebig, Robert Falkenberg, Johannes Pillmann, and Christian Wietfeld. Efficient machine-type communication using multi-metric context-awareness for cars used as mobile sensors in upcoming 5g networks. In *IEEE Vehicular Technology Conference (VTC-Spring)*, Porto, Portugal, June 2018.
- [8] Benjamin Sliwa, Thomas Liebig, Robert Falkenberg, Johannes Pillmann, and Christian Wietfeld. Machine learning based context-predictive car-to-cloud communication using multi-layer connectivity maps for upcoming 5G networks. In *2018 IEEE 88th IEEE Vehicular Technology Conference (VTC-Fall)*, Chicago, USA, August 2018. Accepted for presentation.
- [9] Antonio Viridis, Giovanni Stea, and Giovanni Nardini. *Simulating LTE/LTE-Advanced networks with SimuLTE*, pages 83–105. Springer International Publishing, Cham, 2015.

# Performance Analysis of NB-IoT and eMTC in Challenging Smart City Environments

Pascal Jörke

Lehrstuhl für Kommunikationsnetze  
Technische Universität Dortmund  
pascal.joerke@tu-dortmund.de

The vision of Smart Cities is to enable new use cases including a large number of power restricted Internet of Things devices and challenging channel characteristics. The LTE based communication solutions NB-IoT and eMTC are designed to meet these requirements by integrating new power saving options such as extended Discontinuous Reception and Power Saving Mode as well as the usage of large numbers of repetitions to extend the communication range for coupling losses up to 164 dB. In our work we present a performance analysis of both new technologies and compare data rates, power consumption, latency and spectral efficiency in challenging smart city environments. The results show that, although both technologies use the same power saving techniques as well as repetitions to extend their communication range, both technologies perform different in the context of data size, data rate and coupling loss. While eMTC comes with a 4% better battery lifetime than NB-IoT when considering 144 dB coupling loss, NB-IoT battery lifetime raises to 18% better performance in 164 dB coupling loss scenarios. The overall analysis shows that in coverage areas with a coupling loss of 155 dB or less, eMTC performs better, but requires much more bandwidth. Taking the spectral efficiency into account, NB-IoT is in all evaluated scenarios the better choice and more suitable for future networks with massive number of devices.

The Internet of Things enables many new applications such as Smart Waste Management, Smart Environmental Sensing or Smart Metering. With sensors being placed in

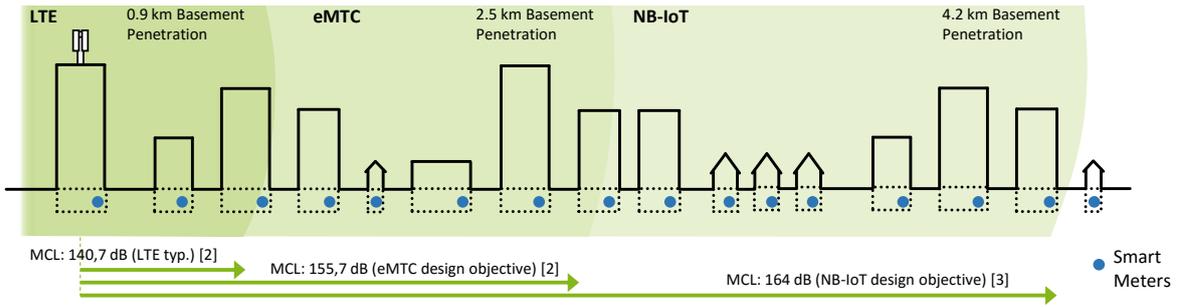


Figure 1: Coverage increase based on 800 MHz Okumura Hata channel models for urban environments and 15 dB additional basement penetration loss [3]

challenging propagation environments such as garages and basements, advanced communication technologies introduce new coverage extension techniques to enable coupling losses up to 164 dB by using mainly repetitions as shown in 1 [4].

Furthermore devices should be able to run over 10 years on a single battery without causing additional costs for maintenance by using extended Discontinuous Reception (eDRX) and a new Power Saving Mode (PSM). While eDRX allows the devices to remain longer in a power saving state between paging occasions, devices in PSM aren't monitoring paging and become unreachable for Mobile Terminated (MT) services, but reduce their power consumption to a bare minimum [2].

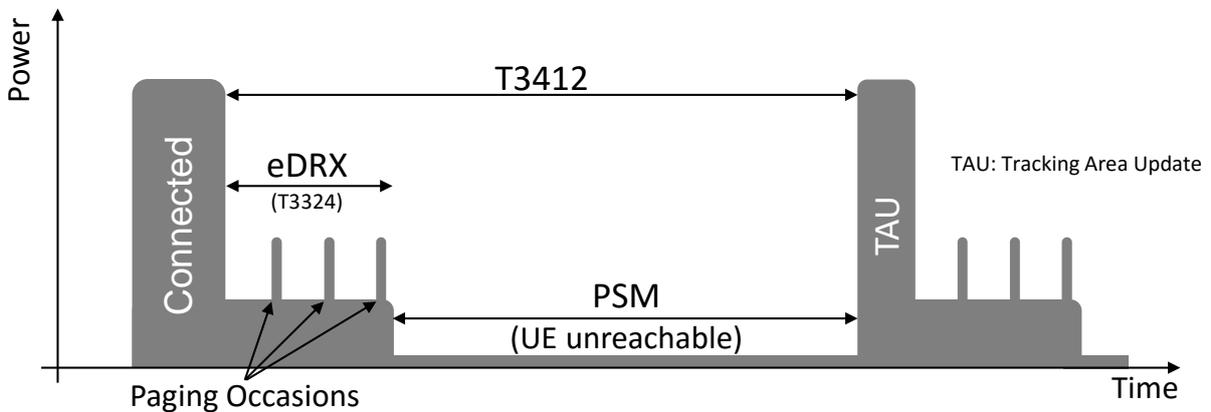


Figure 2: Power reduction in eMTC and NB-IoT by using extended Discontinuous Reception (eDRX) and Power Saving Mode (PSM)

To evaluate the power consumption of NB-IoT and eMTC devices, a new state model is introduced, consisting of five different states:

- *Connected*: After establishing a Radio Resource Control (RRC) connection the device transmits its uplink report.

- *Tail*: After transmitting the uplink report the device keeps active for a time period  $t_{tail}$  in case of possible downlink data. The tail time is defined by the network.
- *eDRX*: The device enters the eDRX mode. It goes in a power saving state and only wakes up in short periods for listening on potential scheduled downlink data.
- *PSM*: After a time  $t_{eDRX}$  defined by T3324 the device enters a deep sleep similar power saving mode where it is not reachable for the network. Because of the power down, the device only consumes several microwatts of power. Although the device chooses the length of T3324 by itself, the network limits the maximum time.
- *TAU*: To signal the network that the device is still alive, it needs to power up after a time  $t_{TAU}$  defined by T3412 to reconnect to the network for a tracking area update and to listen to scheduled downlink data. Although the device chooses the length of T3412, the network limits the maximum time.

This model allows us to compare and evaluate the data rate, power consumption and latency of both technologies in an equal context. The results for 144 dB, 154 dB and 164 dB coupling losses are shown in figure 3 [1].

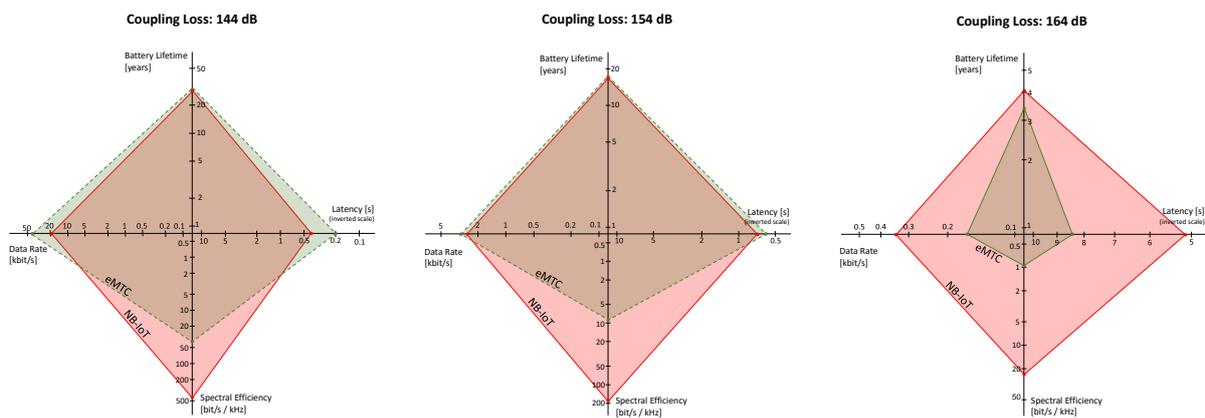


Figure 3: Comparison of NB-LoT and eMTC devices for 84 Bytes uplink data every 24 hours in different coverage conditions

The results show that eMTC performs slightly better in 144 dB and 154 dB coupling loss assumptions than NB-LoT. In the case of 144 dB coupling loss eMTC performs better, providing more than a doubled data rate as well as 4% more battery lifetime than NB-LoT. Additionally the latency of eMTC results in 0.2 s and thus being only half of NB-LoT's latency. However, the extreme propagation conditions of 164 dB coupling loss show a better performance of NB-LoT, which gets along with less repetitions and can thus provide a higher data rate. When taking the spectral efficiency into account, NB-LoT clearly outnumbers eMTC by allocating much less bandwidth and is therefore more suitable for future networks with massive numbers of devices.

This analysis took new Cellular IoT solutions based on LTE into account. In further research the analysis will be extended to license-free spectrum technologies such as LoRaWAN to compare the performance of data rate, latency, battery lifetime and spectral efficiency of licensed spectrum technologies and license-free spectrum technologies. In addition the impact of an increasing number of devices on the battery lifetime, data rate and latency will be evaluated to identify the performance of future highly scaled IoT communication networks.

## References

- [1] Pascal Jörke, Robert Falkenberg, and Christian Wietfeld. Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart City Environments. In *IEEE Global Communications Conference Workshops (GLOBECOM Workshops), Workshop on Green and Sustainable 5G Wireless Networks*, Abu Dhabi, United Arab Emirates, dec 2018.
- [2] O. Liberg, M. Sundberg, E. Wang, J. Bergman, and J. Sachs. *Cellular Internet of Things: Technologies, Standards, and Performance*. Elsevier Science, 2017.
- [3] Stefan Monhof, Stefan Böcker, Janis Tiemann, and Christian Wietfeld. Cellular network coverage analysis and optimization in challenging smart grid environments. In *2018 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, Aalborg, Denmark, oct 2018.
- [4] Sierra Wireless. Coverage Analysis of LTE-M Category-M1. Technical report, Sierra Wireless White Paper, 2017.

# Power Analysis of PhyNode

Mojtaba Masoudinejad  
Lehrstuhl für Förder- und Lagerwesen (FLW)  
Technische Universität Dortmund  
mojtaba.masoudinejad@tu-dortmund.de

Performance of nodes in a wireless sensor network and consequently the overall network relies on balancing the current demand of the hardware and available energy. This work shortly reviews the load of the PhyNode as an embedded entity inside the research lab of PhyNetLab. Overall demand from battery is analyzed for both idle and operational cycle of the PhyNode showing a current deficit in the overall system.

## 1 Introduction

With the rise of Internet of Things (IoT) and Industry 4.0 as its industrial dual, multiple smart small devices are developed to collect data and act decentralized with communication as their core concept. Logistics is not an exception of this revolution with multiple cyber-physical and embedded systems developed for this field [6]. PhyNetLab [2] is a testbed developed for evaluation of different concepts in the field of materials handling and warehousing as two major pillars of the logistics branch. PhyNode is an entity for mounting on transportation boxes used in this field. More than two hundred PhyNodes with different hardware configurations are available in the PhyNetLab. All models of PhyNodes embody an ultra-low power processor, RF communication and some environmental sensors. While availability of energy harvesting module and display is dependent on the model.

Although PhyNetLab has been tested in few applications [2, 4], its optimal operation requires trade-off between service quality and available energy. This shows necessity of power analysis of the PhyNode. Some evaluations such as [1] has been done for individual sections of PhyNode. However, the overall power demand seen from the PhyNode's battery has not been analyzed yet and will be the focus of this work.

## 2 Power module

PhyNode's power module is made of multiple sections. An abstract overview of the complete system including PhotoVoltaic (PV) energy harvesting section is shown in Fig. 1.



Figure 1: Schematic representation of PhyNode's power chain.

An indoor photovoltaic module from *Solems* with 7 cells is used in PhyNode which its behavior and model can be found in [5]. To keep the harvester around its maximum power point while matching its voltage with the battery, a *BQ25505* IC is used. Not only its internal boost converter do the DC-DC matching, but also it keeps the battery within safe operational range. Detailed description of this IC in addition to its behavior model is presented in [3].

PhyNode uses a Li-Polymer battery with a nominal capacity of 1250 mA h, a typical 4.2 V over voltage limit and a cut-off voltage of 3.0 V. Its standard charge and discharge rates are 0.2C. However, measurement at this rate (250 mA) has shown a higher capacity of 1283 mA h. Although this can be considered a marginal capacity difference, it is noticeable for ultra-low power applications such as PhyNode.

A *TPS65290* Power Management IC (PMIC) from Texas Instruments provides different voltage levels from the battery output. In addition to a buck-boost converter, it includes a LDO providing another voltage rail without a switching mechanism. This can be helpful for devices requiring a short term current. While the switched based mechanism requires some switching periods to provide a stable voltage, LDO can directly supply such demand. However, LDOs are generally less efficient due to excess power dissipation through heat.

## 3 Power levels

Measurements in the PhyNetLab has shown that the both extreme cases of photovoltaic generated current will be in the range of 30  $\mu$ A to 140  $\mu$ A. On the other hand, analysis of demands from sub-modules of PhyNode shows that the highest current demand will be during send procedure of the communication module. Considering operational conditions, the highest communication current demand will be about 32 mA. Moreover, processor of the PhyNode requires 100  $\mu$ A per MHz with a maximum of 16 MHz; leading to the highest demand of 1.6 mA. Considering some overhead for other components, a 35 mA is a feasible peak demand.

PMIC acts as an interface between PhyNode's battery and the rest of system. While current demand seen from the battery side will be dependent to its voltage based on the behavior of the PMIC. Therefore, before using the aforementioned limits as a base for system design, real measurements has to be done directly at the output of the battery.

A measurement setup is prepared to replicate the battery with a Source Measurement Unit (SMU). It measures the voltage and current demanded by PMIC while keeping the voltage constant. These devices are connected using a 4-wire connection to enable a non-destructive measurement. A scenario is programmed on a PhyNode which includes all its possible operations; from pushing a button, sending and receiving and even updating the display. Current demands for such operation is measured at different voltage levels starting from 4.2 V with 0.1 V reducing steps. Some examples of these curves for a complete scenario are presented in Fig. 2.

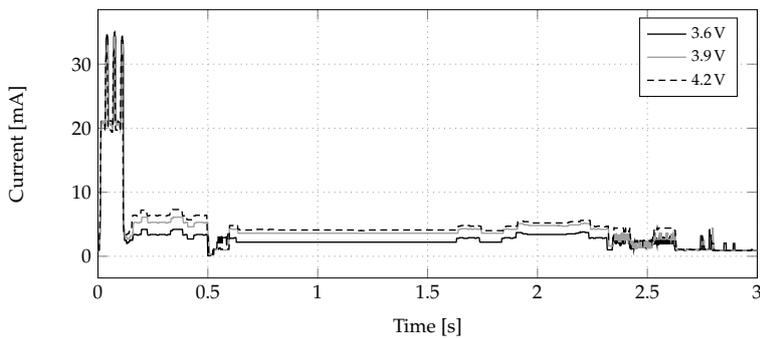


Figure 2: Measured current of PhyNode during a complete cycle at different battery voltages. Peak happens during RF communication send process.

These measurements confirm initial guess of 35 mA as the highest current demanded. Unfortunately, not only reduction of voltage level does not reduce the current demand but also it has negative effect. This is actually a consequence of feeding some parts of the PhyNode with the LDO which simply wastes the excess power. Furthermore, the idle current demand is about 1 mA and is independent from the battery voltage level.

It has been found that reduction of the voltage to the battery minimum is not possible. When the voltage is lower than 3.54 V, PhyNode cannot start all parts of the system during its start check. This effect can also be seen in its current demand shown in Fig. 3.

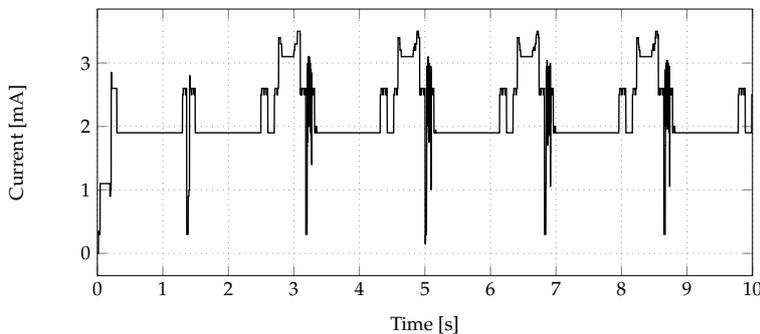


Figure 3: Measured current of PhyNode at 3.53 V causing problem in the startup loop.

## 4 Conclusion

According to measurements, it can be summed up that the PhyNode is only operational in the range of 3.54 V to 4.2 V. Its idle current can be considered independent from the voltage level and around 1 mA while the maximum requested current is about 35 mA. Considering maximum harvested current of 140  $\mu$ A in the PhyNetLab environmental condition this PhyNode version has a current deficit and cannot be considered as an energy-neutral device. Perhaps, software improvements which can reduce the idle power, specifically in the RF module (as the main consumer) can improve this issue. Anyhow, best *idle shelf-time* with the available driver version can be estimation to be around 45 days. In case of operation, this time will be shorten according to the operation profile which its estimation requires further modeling and analysis.

## References

- [1] Markus Buschhoff, Daniel Friesel, and Olaf Spinczyk. Energy Models in the Loop. *Procedia Computer Science*, 130:1063–1068, January 2018.
- [2] R. Falkenberg, M. Masoudinejad, M. Buschhoff, A. K. R. Venkatapathy, D. Friesel, M. ten Hompel, O. Spinczyk, and C. Wietfeld. PhyNetLab: An IoT-based warehouse testbed. In *2017 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 1051–1055, September 2017.
- [3] M. Masoudinejad, M. Magno, L. Benini, and M. ten Hompel. Average Modelling of State-of-the-Art Ultra-low Power Energy Harvesting Converter IC. In *2018 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*, pages 99–104, June 2018.
- [4] M. Masoudinejad, A. K. Ramachandran Venkatapathy, D. Tondorf, D. Heinrich, R. Falkenberg, and M. Buschhoff. Machine Learning Based Indoor Localisation Using Environmental Data in PhyNetLab Warehouse. In *Smart SysTech 2018; European Conference on Smart Objects, Systems and Technologies*, pages 1–8, June 2018.
- [5] Mojtaba Masoudinejad, Manjunath Kamat, Jan Emmerich, Michael ten Hompel, and Saskia Sardesai. A gray box modeling of a photovoltaic cell under low illumination in materials handling application. In *3rd International Renewable and Sustainable Energy Conference (IRSEC)*, pages 1–6, 2015.
- [6] Mojtaba Masoudinejad, Aswin Karthik Ramachandran Venkatapathy, Jan Emmerich, and Andreas Riesner. Smart Sensing Devices for Logistics Application. In *Sensor Systems and Software*, volume 205 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 41–52. Springer International Publishing, Cham, 2017.

# Multiple parallel control of space structures for collaborative manoeuvres

Aswin Karthik Ramachandran Venkatapathy  
Lehrstuhl für Förder- und Lagerwesen  
Technische Universität Dortmund  
aswinkarthik.ramachandran@tu-dortmund.de

The new era of space exploration drives self-assembly and collaborative manoeuvres of space structures in micro-gravity. In this report application of wireless sensor networks in such environments are explored. The concept of TESSERAE is used as one of the target systems for applying low power, low data-rate radio networking in outer space using ultra-low power devices. A collective of wireless networking protocols for decentralized consensus, data replication and networking are presented that can be used for TESSERAE's self-assembly and other space exploration missions. Preparation of the technical demonstration mission in collaboration with Space Exploration Initiative, Media Lab, Massachusetts Institute of Technology is also detailed with the design of hardware and software.

## 1 Introduction

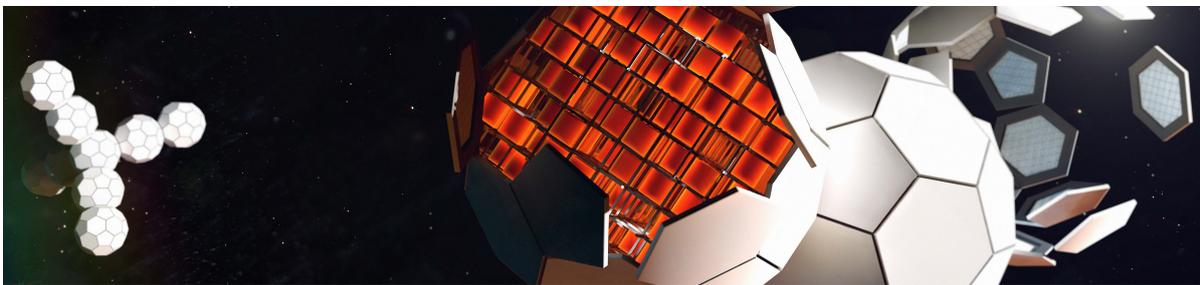


Figure 1: TESSERAE complex assembly to the left with self-assembly of tiles on right.

For space exploration to be scalable and sustainable with increasing human habitats in orbit and on the surface of planets and other habitable celestial bodies. One of the efforts to deploy space habitats was the International Space Station which started in 1998 and it has grown continuously until 2016 to house diverse equipments and experiments with crew. To increase this capability to deploy scalable space structures, it is necessary to break form i.e., the typical cylindrical shape with a radius of the launch vehicle. For such a feat, it is required to engineer self-assembly of structures that can create habitats in space. TESSERAE [3] is one such concept which explores structural formations using tessellated tiles as seen in the  $C_{60}$  Buckminsterfullerene shapes [4]. TESSERAE concept render in Fig. 1 shows the complex structures and self-assembly of the tiles.

TESSERAE concept shows that it is possible to self-assemble tiles with a predetermined trajectory during deployment. For controlling the nature of bonds between the tiles and to manage the state of the tile across all tiles is necessary for control of the tiles. This work adds on to the efforts of sensor nodes on TESSERAE for control [2] and proposes a networking architecture called the decentralized brain [4] where data replication is performed to replicate the states of a tile across all tiles.

## 2 Networking Requirements

For self-assembly in space, the networking application should provide ad hoc capabilities. When a node is faulty or leaves the network it must be mitigated by self-healing capabilities. For synchronous behaviour of systems to perform actions together and to interact with each other in space, the requirements for the network should also consider a reliable data dissemination method for network-wide agreement and global configuration management in a decentralized manner. In the case of TESSERAE, swarm behaviour is transient during assembly. Therefore, nodes require low latency data dissemination as well as adaptive replication strategies during interaction to enable precise control of interacting elements or guarantees in information propagation. In addition to the TESSERAE self-assembly swarm, we note an applicability for this work to other space applications, such as a space-fed phased array antenna swarm. Here, each element of the array is free to roam within a certain area, while compensating for position error via selectively delaying and amplifying signals to generate a coherent radio signal. The communication architecture we are proposing would facilitate relative position tracking between swarm elements. The communication architecture should facilitate relative position tracking between swarm elements and the necessary variable-delay updating and resynchronisation of timing between emitters to achieve coherently cooperating RF elements. This is in contrast to a fixed-array antenna, where elements are held rigidly along a plane and can use a standard delay based on fixed position along the antenna element line.

## 2.1 Decentralized brain

In *distributed computing*, state machine replication or data structure replication is a common method for implementing a fault-tolerant service by replicating servers and coordinating client interactions with server replicas. The fundamental communication primitive that is used in distributed computing to achieve consensus between processes is atomic broadcast. An atomic broadcast is where all correct nodes in a network receive the same set of messages in the same order. Atomic broadcasts are an important distributed computing primitive for consensus which is achieved by Glossy, VHT, and LoWPAN to uniquely integrate these into a robust, low-power, low data-rate distributed network for a decentralized brain [4].

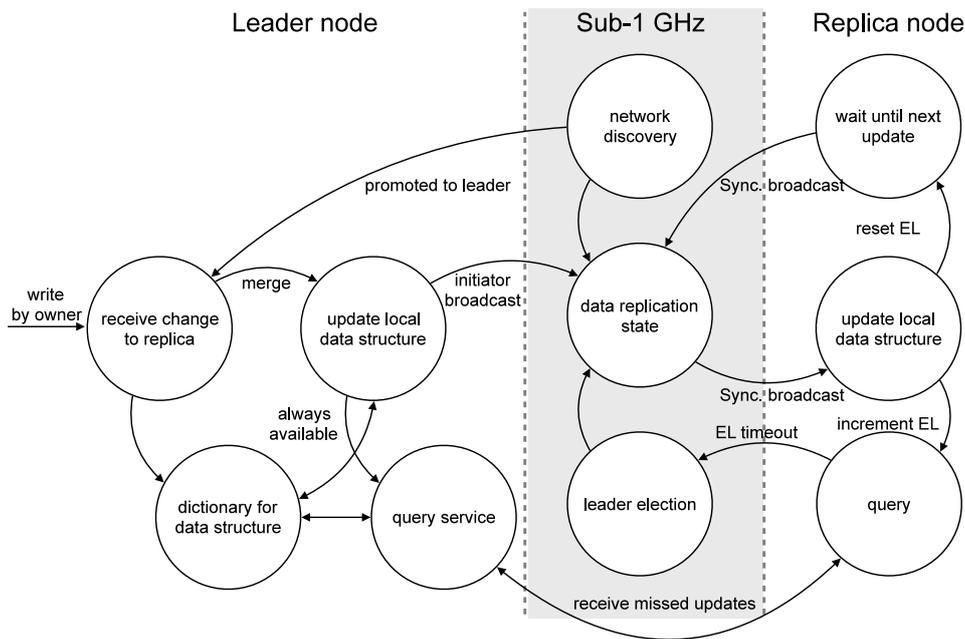


Figure 2: TESSERAE complex assembly to the left with self-assembly of tiles on right.

In figure 2.1, the states for a dual band, decentralized leader based data replication protocol for ultra-low power devices is shown. The data replication is a modular communication primitive that provides easier state management between nodes. The time required per update propagation  $t_p$  can be estimated using the sum of individual timing requirements in the networking architecture. As per the requirements of TESSERAE, required broadcast is 22 bytes sensor data transfer; here a reliable data-rate of 124 bytes per replication round, excluding the physical layer header and the replication header, is available. The data-rate is increased by decoupling the meta-data and serialization information for data replication using the message type bit in the CMD field; this maximizes the data replication since the same data structure is persisted with consistency more frequently than the changes to the information itself in collaborative control maneuvers. In large-scale, low-power networks, the decentralized brain concept is achieved using data replication

with variably participating nodes, controlled by the *cluster differentiate setting* in the CMD field of the payload. A new *brain* is initialized using the lowpan network among the nodes that need to perform control in precise manner, thus decreasing the number of hops and also reducing the latency while propagating the states to the participating nodes in <3 ms. Cluster differentiate setting is designed to handle up to 16 different leaders where the nodes can choose to participate in data replication. The rate at which data is replicated can be set up by the leader node, allowing for use dependent frequency in data replication (i.e., the whole system is updated with its slow-changing state of assembly less-frequently than the updates required for quick, precise control between the interacting entities). This provides for a guarantee in communication, which is critical for precise control, as well as modular communication such as point-to-point.

### 3 Future work

The concept of decentralized brain with hardware design and development of the sensor nodes of TESSERAE were developed for a technical demonstration mission which will be tested in micro-gravity. Future work will include large scale performance tests for the concept of decentralized brain [4] along with the context broker relaying information to the higher layers of the control system. These performed tests will be evaluated in context with two space experiments one of which will be TESSERAE. A large scale sensor floor implementation will be used for testing and defining a networking stack for the decentralized brain with open source implementation such as Contiki-NG.

### References

- [1] Ariel Ekblaw and Joseph Paradiso. Self-assembling Space Structures: Buckminsterfullerene Sensor Nodes. In *2018 AIAA/AHS Adaptive Structures Conference*, Kissimmee, Florida, January 2018. American Institute of Aeronautics and Astronautics.
- [2] Ariel Ekblaw and Joseph A. Paradiso. TESSERAE : Self-assembling shell structures for space exploration. In *Proceedings of the IASS Symposium 2018 Creativity in Structural Design*, Boston, Massachusetts, July 2018.
- [3] Aswin Karthik Ramachandran Venkatapathy, Ariel Ekblaw, Michael ten Hompel, and Joseph Paradiso. Decentralized brain in low data-rate, low power networks for collaborative manoeuvres in space. In *The 6th Annual IEEE International Conference on Wireless for Space and Extreme Environments*, Huntsville, Alabama, November 2018.

# Reliable Wireless Positioning based on IEEE 802.15.4a Ultra-Wideband for Low-Power Applications

Janis Tiemann

Lehrstuhl für Kommunikationsnetze  
Technische Universität Dortmund  
janis.tiemann@tu-dortmund.de

Recent developments in wireless positioning based on IEEE 802.15.4a gained broad attention for various applications. However, in most current research many aspects required for low-power applications are not considered. Therefore, this work aims on overcoming the limitations of most schemes, by focusing on the energy efficiency of such systems. Although recent advances in scalable wireless positioning topologies pave the way for improved energy efficiency, many of these approaches do not consider large channel utilization. This work aims to provide guaranteed update rates, reliable error-mitigating localization results and low-energy consumption.

## 1 Introduction

Wireless localization using ultra-wideband (UWB) communication technology enabled a wide field of research in recent years. Previous work considered the use of this technology as a control feedback for autonomous robotic systems such as indoor-navigating unmanned aerial vehicles (UAVs), see [1]. Further research evaluated the use for precision alignment of electric vehicles for wireless power transfer (WPT), see [2]. Based on those developments an open-source solution for scalable time-difference of arrival (TDOA) based localization emerged in previous work, see [3]. It could be shown, that advanced applications are feasible using this approach, see [4]. Recent work extended the usable range of the UAVs through monocular vision based simultaneous localization

and mapping (SLAM) augmentation of the UWB system, see [5]. In other work, it could be shown, that reliable localization results with the low-power methods developed in this work are feasible, see [6]. Participation in the EvAAL'16 competition could validate the practical feasibility, see [7].

Although accuracy and reliability are core aspects of wireless localization systems, this work's main focus lies within the resource efficiency in providing quality of service (QoS) guaranteed localization specific wireless channel access. In contrast to many existing localization topologies that utilize two-way ranging (TWR) based message exchanges or random access based uncoordinated TDOA, this work proposes to use centralized coordination of the mobile units to achieve time-division multiple access (TDMA). However, coordination requires the use of additional resources. Our goal is to quantify this overhead and provide guidelines such that wireless localization system designers are enabled to base their system design decisions on the given analysis.

## 2 Random vs. Coordinated Access for Low-Power TDOA-based Ultra-Wideband Localization

In the experimental analysis in [8] it could be shown that random access packet-based access in the IEEE 802.15.4a UWB PHY has special properties through the sparse burst-

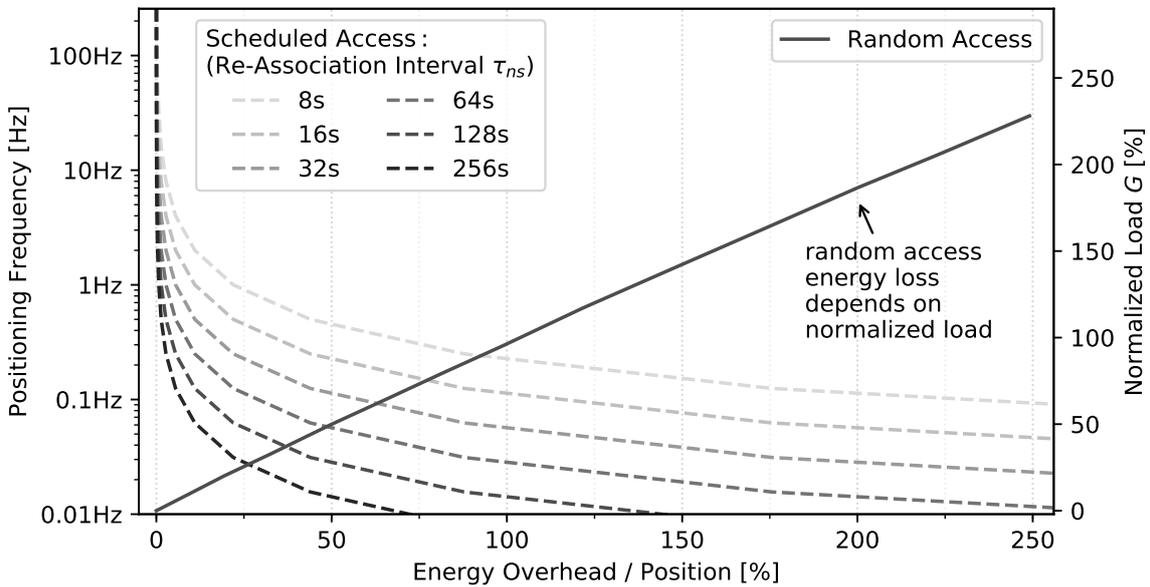


Figure 1: Analysis of the energy loss over the positioning frequency under different system loads for the random access scheme (R-TDOA) compared to the proposed scheduled access scheme under different positioning frequencies.

position modulation binary phase shift keying (BPM-BPSK) scheme. The PHY requires tight time-synchronization for demodulation. Therefore, interfering access does not necessarily cause destructive collisions. Depending on the channel configuration, the overall throughput  $S_{p:rtdoa}$  can be approximated as almost collision-free, see (1).

$$S_{p:rtdoa} \approx (1 - e^{-G}) \quad (1)$$

The rate of successfully obtained positions  $R_p$  divided by the rate capacity of the channel  $R_m$  defined by the duration  $T_{PPDU}$  of the physical protocol data unit (PPDU) leads to the normalized localization throughput  $S_p$ . Hence, the load induced traffic  $G$  is based as the actual rate of frames on the channel  $R_f$ , see (2).

$$R_m = 1/T_{PPDU} \quad S_p = R_p/R_m \quad G = R_f/R_m \quad (2)$$

For the traffic depended analysis the number of mobile nodes (tags) is considered to be large  $N_u \gg 1$ . The energy overhead per successfully obtained position estimation is depicted in Fig. 1. Here, it is clearly visible, that for random access, the energy loss directly corresponds with the overall normalized system load  $G$ . For the scheduled channel access, the re-association interval  $\tau_{ns}$ , which defines how often the mobile unit needs to listen to a central coordinator defines the energy overhead per position. Here, the positioning frequency is the main factor that defines the ratio of successfully obtained positions vs. the amount of re-association and synchronization overhead. It should be noted, that based on the technical specifications of the currently most frequently used transceiver chipsets, the reception is assumed to be twice as expensive as transmission.

It is clearly visible, that for low-frequency localization using coordinated access, high re-association intervals are required for it to be more energy efficient than random access. The re-association interval is mainly influenced by the dynamics and quality of the mobile clock sources. When, higher dynamics are required such as in the tracking of persons or robots, coordinated access performs almost always better considering a typical overall normalized load.

### 3 Conclusion and Further Research

In this work, the trade-off between random and coordinated access in TDOA-based ultra-wideband localization is analyzed. Analytical models for the energy overhead, based on a set of system properties are given. Based on the models, general guidelines for system design considerations could be derived. Future work will focus on providing a low-power system integrated implementation of the algorithms for coordinated access.

## References

- [1] J. Tiemann, F. Schweikowski, and C. Wietfeld. Design of an UWB indoor-positioning system for UAV navigation in GNSS-denied environments. In *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*, Oct 2015.
- [2] J. Tiemann, J. Pillmann, S. Böcker, and C. Wietfeld. Ultra-wideband aided precision parking for wireless power transfer to electric vehicles in real life scenarios. In *IEEE Vehicular Technology Conference (VTC-Fall)*, Montréal, Canada, Sep 2016.
- [3] J. Tiemann, F. Eckermann, S. Böcker, and C. Wietfeld. ATLAS - an open-source TDOA-based ultra-wideband localization system. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Alcalá de Henares, Madrid, Spain, Oct 2016.
- [4] J. Tiemann and C. Wietfeld. Scalable and precise multi-UAV indoor navigation using TDOA-based UWB localization. In *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Sapporo, Japan, (Awarded Best Paper) Sep 2017.
- [5] J. Tiemann, A. Ramsey, and C. Wietfeld. Enhanced UAV indoor navigation through SLAM-augmented UWB localization. In *IEEE International Conference on Communications Workshops (ICC Workshops): The IEEE International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (CCNCPS)*, Kansas City, USA, May 2018.
- [6] J. Tiemann, L. Koring, P. Gorczak, and C. Wietfeld. Improving the robustness of control-grade ultra-wideband localization. In *3rd IFAC Conference on Embedded Systems, Computer Intelligence and Telematics (CESCIT 2018)*, Faro, Portugal, Jun 2018.
- [7] M. D. C. Pérez Rubio, C. Losada-Gutiérrez, F. Espinosa, J. Macias-Guarasa, J. Tiemann, F. Eckermann, C. Wietfeld, M. Katkov, S. Huba, J. Ureña, J. M. Villadangos, D. Gualda, E. Díaz, R. Nieto, E. Santiso, P. D. Portillo, and M. Martínez. A realistic evaluation of indoor robot position tracking systems: The IPIN 2016 competition experience. *Measurement*, Nov 2018.
- [8] J. Tiemann, F. Eckermann, S. Böcker, and C. Wietfeld. Multi-user interference and wireless clock synchronization in TDOA-based UWB localization. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Alcalá de Henares, Madrid, Spain, Oct 2016.





Subproject A6  
Resource-efficient Graph Mining

Nils Kriege      Petra Mutzel  
Frank Weichert

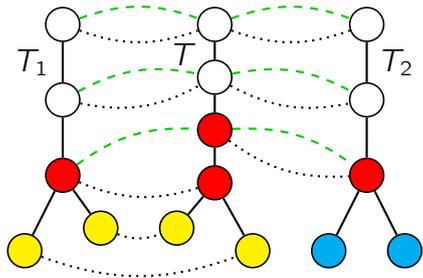
# Largest Weight Common Subtree Embeddings with Distance Penalties

Andre Droschinsky  
Chair of Algorithm Engineering (LS11)  
TU Dortmund  
andre.droschinsky@tu-dortmund.de

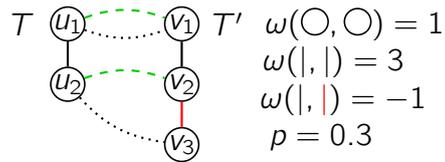
The largest common embeddable subtree problem asks for the largest possible tree embeddable into two input trees and generalizes the classical maximum common subtree problem. Several variants of the problem in labeled and unlabeled rooted trees have been studied, e.g., for the comparison of evolutionary trees. We consider a generalization, where the sought embedding is maximal with regard to a weight function on pairs of labels. We support rooted and unrooted trees with vertex and edge labels as well as distance penalties for skipping vertices. This variant is important for many applications such as the comparison of chemical structures and evolutionary trees. Our algorithm computes the solution from a series of bipartite matching instances, which are solved efficiently by exploiting their structural relation and imbalance. Our analysis shows that our approach improves or matches the running time of the formally best algorithms for several problem variants. The results were published at MFCS 2018 [4].

## 1 Largest Weight Common Subtree Embeddings

Various variants for comparing trees have been proposed and investigated [10]. Most of them assume rooted trees, which may be ordered or unordered. Algorithms tailored to the comparison of evolutionary trees typically assume only the leaves to be labeled, while others support labels on all vertices or do not consider labels at all. The well-known agreement subtree problem, for example, considers the case, where only the leaf nodes are labeled, with no label appearing more than once per tree [7]. We discuss the



(a) Labeled MCS (green dashed lines) and LaCSE (black dotted lines) between  $T$  and  $T_i$ ,  $i \in \{1, 2\}$ .



(b) The black embedding has weight 1.7, since the vertex  $v_2$  is skipped and therefore the penalty  $p$  is applied; the weight between the edges is not added. The green embedding has weight 5; 2 from the vertices, 3 from the path  $(u_1, u_2)$  mapped to  $(v_1, v_2)$ .

Figure 1: **a)** Although 'intuitively'  $T$  is more similar to  $T_1$  than to  $T_2$ , both MCSs have size 3. However, the LaCSE between  $T$  and  $T_1$  has 6 mapped vertices. **b)** Two weighted embeddings; one with a skipped vertex, the other where the edge labels contribute to the weight.

approaches most relevant for our work. Gupta and Nishimura [5] investigated the *largest common embeddable subtree* problem in unlabeled rooted trees. Their definition is based on topological embedding (or homeomorphism) and allows to map edges of the common subtree to vertex-disjoint paths in the input trees. The algorithm uses the classical idea to decompose the problem into subproblems for smaller trees, which are solved via bipartite matching. A solution for two rooted trees with at most  $n$  vertices is computed in time  $\mathcal{O}(n^{2.5} \log n)$  using a dynamic programming approach. Many algorithms do not support trees, where leaves and the inner vertices both have labels. A notable exception is the approach by Kao et al. [6], where only vertices with the same label may be mapped (LaCSE). This algorithm generalizes the approach by Gupta/Nishimura and improves its running time to  $\mathcal{O}(\sqrt{d}D \log \frac{2n}{d})$ , where  $D$  denotes the number of vertex pairs with the same label and  $d$  the maximum degree of all vertices. In our work we consider the problem of finding a *largest weight common subtree embedding* (LaWeCSE) between unrooted trees, where matching vertices are not required to have the same label, but their degree of agreement is determined by a weight function. To prevent arbitrarily long paths which are mapped to a common edge we study a linear distance penalty for paths of length greater than 1. Note that, by choosing a high distance penalty, we solve the maximum common subtree (MCS) problem as a special case.

We propose and analyze algorithms for finding largest weight common subtree embeddings. Our method requires to solve a series of bipartite matching instances as subproblem, which dominates the total running time. We build on recent results by Ramshaw and Tarjan [8, 9] for unbalanced matchings. Let  $T$  and  $T'$  be labeled rooted trees with  $k := |T|$  and  $l := |T'|$  vertices, respectively, and  $\Delta := \min\{\Delta(T), \Delta(T')\}$  the smaller degree of the two input trees. For real-valued weight functions we prove a time bound of  $\mathcal{O}(kl\Delta)$ . For integral weights bounded by a constant  $C$  we prove a running time of

$\mathcal{O}(kl\sqrt{\Delta}\log(\min\{k, l\}C))$ . This is an improvement over the algorithm by Kao et al. [6] if there are only few labels and the maximum degree of one tree is much smaller than the maximum degree of the other. In addition, we support weights and a linear penalty for skipped vertices. The weight  $\mathcal{W}(\varphi)$  of a common subtree embedding  $\varphi$  is the sum of the weights  $\omega(u, \varphi(u))$  of all vertices  $u$  mapped by  $\varphi$  plus the weights  $\omega_p(P, \varphi(P))$  of all mapped paths  $P$ . If both  $P$  and  $\varphi(P)$  are single edges,  $\omega_p(P, \varphi(P))$  is determined by the two edges (e. g. their labels); otherwise  $\omega_p(P, \varphi(P)) = -ip$ , where  $i$  is the number of inner vertices and  $p \geq 0$  is the distance penalty.

Moreover, the algorithm by Kao et al. [6] is designed for rooted trees only. A straight forward approach to solve the problem for unrooted trees is to try out all pairs of possible roots, which results in an additional  $\mathcal{O}(kl)$  factor. However, our algorithm exploits the fact that there are many similar matching instances using techniques related to [1, 3]. This includes computing additional matchings of cardinality two. For unrooted trees and real-valued weight functions we prove the same  $\mathcal{O}(kl\Delta)$  time bound as for rooted trees. This leads to an improvement over the formally best algorithm for solving the maximum common subtree problem, for which a time bound of  $\mathcal{O}(kl(\Delta + \log d))$  has been proven [3]. Fig. 1 exemplifies LaCSE in comparison to MCS as well as embeddings with skipped vertices. We proved the following results.

**Theorem 1** *Let  $T$  and  $T'$  be rooted vertex and/or edge labeled trees. Let  $w$  be a weight function,  $\Delta = \min\{\Delta(T), \Delta(T')\}$ , and  $p$  be a distance penalty.*

- *A LaWeCSE between  $T$  and  $T'$  can be computed in time  $\mathcal{O}(|T||T'|\Delta)$  and space  $\mathcal{O}(|T||T'|)$ .*
- *If the weights are integral and bounded by a constant  $C$ , a LaWeCSE can be computed in time  $\mathcal{O}(|T||T'|\sqrt{\Delta}\log(C\min\{|T|, |T'|\}))$ .*

**Theorem 2** *Let  $T$  and  $T'$  be (unrooted) vertex and/or edge labeled trees. Let  $w$  be a weight function,  $\Delta = \min\{\Delta(T), \Delta(T')\}$ , and  $p$  be a distance penalty. A LaWeCSE <sub>$u$</sub>  between  $T$  and  $T'$  can be computed in time  $\mathcal{O}(|T||T'|\Delta)$  and space  $\mathcal{O}(|T||T'|)$ .*

## 2 Future work

We integrated the technique to skip vertices into our software to compare outerplanar graphs using the block and bridge preserving MCS [2]. First experimental results on molecular graphs are promising (most molecular graphs are outerplanar), since sometimes the difference between two molecules is just the length of a chain connecting isomorphic parts of the molecular graphs. We plan to further investigate the practical usefulness. We also want to include bioisosteres into the software as well as improve the running time for non outerplanar graphs.

## References

- [1] Moon Jung Chung.  $O(n^{2.5})$  time algorithms for the subgraph homeomorphism problem on trees. *Journal of Algorithms*, 8(1):106–112, 1987.
- [2] Andre Droschinsky, Nils Kriege, and Petra Mutzel. *Finding Largest Common Substructures of Molecules in Quadratic Time*, pages 309–321. Springer International Publishing, Cham, 2017.
- [3] Andre Droschinsky, Nils M. Kriege, and Petra Mutzel. Faster Algorithms for the Maximum Common Subtree Isomorphism Problem. In Piotr Faliszewski, Anca Muscholl, and Rolf Niedermeier, editors, *41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016)*, volume 58 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [4] Andre Droschinsky, Nils M. Kriege, and Petra Mutzel. Largest Weight Common Subtree Embeddings with Distance Penalties. In Igor Potapov, Paul Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018)*, volume 117 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 54:1–54:15, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [5] Arvind Gupta and Naomi Nishimura. Finding largest subtrees and smallest supertrees. *Algorithmica*, 21:183–210, 1998.
- [6] Ming-Yang Kao, Tak-Wah Lam, Wing-Kin Sung, and Hing-Fung Ting. An even faster and more unifying algorithm for comparing trees via unbalanced bipartite matchings. *Journal of Algorithms*, 40(2):212–233, 2001.
- [7] Daniel M. Martin and Bhalchandra D. Thatte. The maximum agreement subtree problem. *Discrete Applied Mathematics*, 161(13-14):1805–1817, 2013.
- [8] L. Ramshaw and R. E. Tarjan. A weight-scaling algorithm for min-cost imperfect matchings in bipartite graphs. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 581–590, Oct 2012.
- [9] Lyle Ramshaw and Robert Tarjan. On minimum-cost assignments in unbalanced bipartite graphs, 04 2012.
- [10] Gabriel Valiente. *Algorithms on Trees and Graphs*. Springer-Verlag, Berlin, 2002.

# Learning on Irregularly Structured Data and its Applications

Matthias Fey

Department of Computer Graphics

TU Dortmund University

matthias.fey@tu-dortmund.de

Graph neural networks (GNNs) have been recently emerged as one of the most successful approaches to tackle the challenging and highly active field of representation learning on irregular domains like graphs and point clouds for tasks such as classification or segmentation [1]. With the generalization of the convolutional and pooling layers to these domains, GNNs are able to (hierarchically) extract localized embeddings by passing, transforming, and aggregating information, significantly advancing the state-of-the-art on many challenging benchmark tasks [1, 3]. Here, we introduced a novel B-spline based convolutional operator to apply CNNs on irregularly structured data and further demonstrated its applications in two follow-up works. We also presented the PyTorch Geometric framework that allows for fast graph representation learning by leveraging GPU capabilities, with many use cases in research and production.

## 1 Spline-based Convolutional Operator

Our *spline-based convolutional operator* (*SplineConv*) [2] is a generalization of the traditional convolutional operator for irregularly structured data. Let  $\mathcal{G} = (\mathbf{X}, \mathbf{E})$  be a directed graph represented by its node feature matrix  $\mathbf{X} \in \mathbb{R}^{N \times M}$  with  $N$  nodes and  $M$  features each, and its (sparse) adjacency tensor  $\mathbf{E}^{N \times N \times D}$  with  $D$ -dimensional edge features. Let  $\mathcal{N}(i)$  further denote the neighborhood set of node  $i$ .

Our proposed convolutional operator aggregates node features in local neighborhoods weighted by a trainable, continuous kernel function. The node features  $\mathbf{x}_i$  represent

features on an irregular structure, whose relations are locally defined by its edge features  $\mathbf{e}_{i,j}$ . Therefore, when locally aggregating feature values in a node’s neighborhood, edge features are used to determine *how* the features are aggregated and the content of the node features define *what* is aggregated.

Therefore, convolution over neighboring features for a vertex  $i$  is defined by

$$(\mathbf{X} \star \mathbf{g})_i = \frac{1}{\mathcal{N}(i)} \sum_{m=1}^M \sum_{j \in \mathcal{N}(i)} x_{j,m} \cdot g_m(\mathbf{e}_{i,j}) \quad (1)$$

where  $\mathbf{g} = (g_1, \dots, g_M)$  defines  $M$  continuous kernel functions, which weight the components of  $\mathbf{x}_j$  based on the local relation  $\mathbf{e}_{i,j}$  between vertex  $i$  and  $j$ .

A kernel function  $g_m$  is parameterized by a fixed number of trainable parameters. For computing  $g_m(\mathbf{e}_{i,j})$ , the kernel function relies on the product of B-spline basis functions of a user-defined degree and can thus be evaluated very efficiently due to the local support property of B-splines. Using this operator, we can construct a convolutional layer that can be used in deep neural network architectures to train models on (embedded) graph data.

## 2 Applications

We demonstrated possible applications of SplineCNN in two follow-up works.

**Recognizing Cuneiform Signs.** The cuneiform script constitutes one of the earliest systems of writing and is realized by wedge-shaped marks on clay tablets. A tremendous number of cuneiform tablets have already been discovered and are incrementally digitalized and made available to automated processing.

As reading cuneiform script is still a manual task, we addressed the real-world application of recognizing cuneiform signs using our spline-based convolutional operator [4]. Due to limited available annotated real-world training data, we strengthened the performance of our CNN by enriching the training set by augmented examples. By applying random affine transformations like translation, scaling and rotation on the cuneiform’s graph representation (*cf.* Figure 1), we could further improve our model performance from 87.37% to 93.53%.

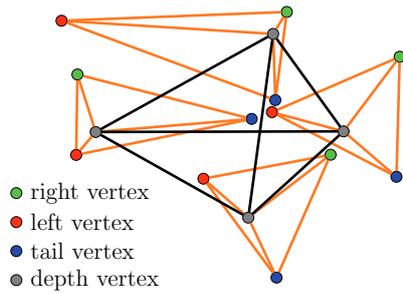


Figure 1: Cuneiform graph representation.

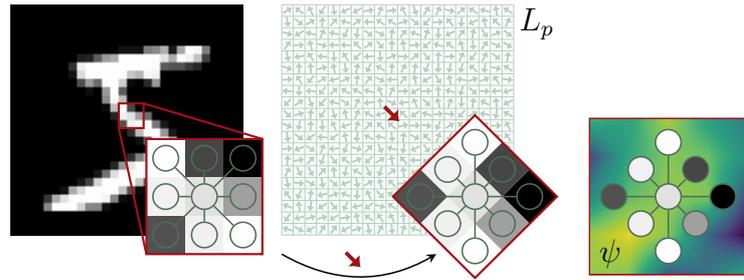


Figure 2: Locally rotating receptive fields based on equivariant pose vectors.

**Sparse group convolution.** Our work on *group equivariant capsule networks* [5] introduced a framework that brings guaranteed equivariance and invariance properties to the capsule network idea [6]. Here, we presented a generic *routing by agreement* algorithm defined on elements of a Lie group, *e.g.*,  $SO(2)$ , and proved that equivariance of output pose vectors, as well as invariance of output activations hold under certain conditions. We further showed that the obtained equivariant pose vectors can be further leveraged to build convolutional neural network architectures that inherit those properties. Calculation of the convolutions can be performed by applying the inverse transformation to the local input using the capsule’s pose vector (*cf.* Figure 2). By interpreting images as grid-graphs, we used our spline-based convolution to operate on locally transformed receptive fields.

### 3 Framework: PyTorch Geometric

We also introduced *PyTorch Geometric*<sup>1</sup>, a geometric deep learning framework [1] for PyTorch, which consists of a variety of published methods for learning on irregularly structured input data such as graphs or point clouds. PyTorch Geometric achieves high data throughput by leveraging sparse GPU acceleration, by providing custom CUDA kernels and by introducing efficient mini-batch handling for examples with potentially different size.

**Convolutional Operators.** In practice, almost all convolutional operators on irregularly structured data can be implemented by gathering and scattering of node features as well as making use of broadcasting for element-wise kernel evaluation. Although working on irregularly structured input, this scheme can be heavily accelerated by the GPU. We already implemented over 10 convolutional operators often found in literature [3].

<sup>1</sup>GitHub repository: [https://github.com/rusty1s/pytorch\\_geometric](https://github.com/rusty1s/pytorch_geometric)

**Mini-batch Handling.** Our framework supports batches of multiple graph instances (of potentially different size) by automatically creating a single (sparse) block-diagonal adjacency matrix and concatenating feature matrices in the node dimension. In this way, neighborhood aggregation methods can be applied normally, since no messages are exchanged between disconnected subgraphs.

**Datasets.** We currently support over 70 graph and point cloud benchmark datasets<sup>2</sup>. All provided datasets are automatically downloaded and processed into a consistent data format. In addition, users can create their own datasets by following simple guidelines.

## 4 Future Work

GNNs hold the potential to be applied to a large number of different datasets and deep learning problems. We plan to further develop new methods and transfer methods from known traditional CNNs for irregularly structured data, e.g. pooling operators which, conversely, have not yet received a great amount of attention.

## References

- [1] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. In *IEEE Signal Processing Magazine*, 2017.
- [2] M. Fey, J. E. Lenssen, F. Weichert, and H. Müller. SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In *CVPR*, 2018.
- [3] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017.
- [4] N. M. Kriege, M. Fey, D. Fisseler, P. Mutzel, and F. Weichert. Recognizing cuneiform signs using graph based methods. In *PMLR*, 2018.
- [5] J. E. Lenssen, M. Fey, and P. Libuschewski. Group equivariant capsule networks. In *NeurIPS*, 2018.
- [6] S. Sabour, N. Frosst, and G. E. Hinton. Dynamic routing between capsules. In *NIPS*, 2017.

---

<sup>2</sup>For example, every graph kernel benchmark dataset from <http://graphkernels.cs.tu-dortmund.de> can be used.

# Weisfeiler and Leman go neural: Higher-order graph embeddings

Christopher Morris  
Chair of Algorithm Engineering  
TU Dortmund University  
christopher.morris@tu-dortmund.de

We report about our recent progress in relating graph kernels and graph neural networks: In recent years, graph neural networks (GNNs) have emerged as a powerful neural architecture to learn vector representations of nodes and graphs in a supervised, end-to-end fashion. Up to now, GNNs have only been evaluated empirically—showing promising results. The following work investigates GNNs from a theoretical point of view and relates them to the 1-dimensional Weisfeiler-Leman graph isomorphism heuristic (1-WL). We show that GNNs have the same expressiveness as the 1-WL in terms of distinguishing non-isomorphic (sub-)graphs. Hence, both algorithms also have the same shortcomings. Based on this, we propose a generalization of GNNs, so-called  $k$ -dimensional GNNs ( $k$ -GNNs), which can take higher-order graph structures at multiple scales into account. These higher-order structures play an essential role in the characterization of social networks and molecule graphs. Our experimental evaluation confirms our theoretical findings as well as confirms that higher-order information is useful in the task of graph classification and regression. This work has been accepted at the AAAI Conference 2019.

## 1 The Weisfeiler-Leman algorithm and GNNs

Graph-structured data is ubiquitous across application domains ranging from chemo- and bioinformatics to image and social network analysis. To develop successful machine learning models in these domains, we need techniques that can exploit the rich information inherent in graph structure, as well as the feature information contained within a graph's

nodes and edges. In recent years, numerous approaches have been proposed for machine learning graphs—most notably, approaches based on graph kernels [11] or, alternatively, using graph neural network algorithms [5].

Kernel approaches typically fix a set of features in advance—e.g., indicator features over subgraph structures or features of local node neighborhoods. For example, one of the most successful kernel approaches, the *Weisfeiler-Lehman subtree kernel* [10], which is based on the 1-dimensional Weisfeiler-Leman graph isomorphism heuristic [3, pp. 79 ff.], generates node features through an iterative relabeling, or *coloring*, scheme: First, all nodes are assigned a common initial color; the algorithm then iteratively recolors a node by aggregating over the multiset of colors in its neighborhood, and the final feature representation of a graph is the histogram of the resulting node colors. By iteratively aggregating over local node neighborhoods in this way, the WL subtree kernel is able to effectively summarize the neighborhood substructures present in a graph. However, while powerful, the WL subtree kernel—like other kernel methods—is limited because this feature construction scheme is fixed (i.e., it does not adapt to the given data distribution). Moreover, this approach—like the majority of kernel methods—focuses only on the graph structure and cannot interpret continuous node and edge labels, such as real-valued vectors which play an important role in applications such as bio- and chemoinformatics.

Graph neural networks (GNNs) have emerged as a machine learning framework addressing the above challenges. Standard GNNs can be viewed as a neural version of the 1-WL algorithm, where colors are replaced by continuous feature vectors and neural networks are used to aggregate over node neighborhoods [4, 7]. In effect, the GNN framework can be viewed as implementing a continuous form of graph-based “message passing”, where local neighborhood information is aggregated and passed on to the neighbors [2]. Let  $(G, l)$  be a labeled graph with an initial node coloring  $f^{(0)} : V(G) \rightarrow \mathbb{R}^{1 \times d}$  that is *consistent* with  $l$ . This means that each node  $v$  is annotated with a feature  $f^{(0)}(v)$  in  $\mathbb{R}^{1 \times d}$  such that  $f^{(0)}(u) = f^{(0)}(v)$  if and only if  $l(u) = l(v)$ . Alternatively,  $f^{(0)}(v)$  can be an arbitrary real-valued feature vector associated with  $v$ . Examples include continuous atomic properties in chemoinformatic applications where nodes correspond to atoms, or vector representations of text in social network applications. A GNN model consists of a stack of neural network layers, where each layer aggregates local neighborhood information, i.e., features of neighbors, around each node and then passes this aggregated information on to the next layer. A basic GNN model can be implemented as follows [5]. In each layer  $t > 0$ , we compute a new feature

$$f^{(t)}(v) = \sigma \left( f^{(t-1)}(v) \cdot W_1^{(t)} + \sum_{w \in N(v)} f^{(t-1)}(w) \cdot W_2^{(t)} \right) \quad (1)$$

in  $\mathbb{R}^{1 \times e}$  for  $v$ , where  $W_1^{(t)}$  and  $W_2^{(t)}$  are parameter matrices from  $\mathbb{R}^{d \times e}$ , and  $\sigma$  denotes a component-wise non-linear function, e.g., a sigmoid or a ReLU.<sup>1</sup>

---

<sup>1</sup>For clarity of presentation we omit biases.

Both  $f_{\text{aggr}}^{W_1}$  and  $f_{\text{merge}}^{W_2}$  may be arbitrary differentiable, permutation-invariant functions (e.g., neural networks), and, by analogy to Equation 1, we denote their parameters as  $W_1$  and  $W_2$ , respectively. By deploying a trainable neural network to aggregate information in local node neighborhoods, GNNs can be trained in an end-to-end fashion together with the parameters of the classification or regression algorithm, possibly allowing for greater adaptability and better generalization compared to the kernel counterpart of the classical 1-WL algorithm.

Up to now, the evaluation and analysis of GNNs has been largely empirical, showing promising results compared to kernel approaches, see, e.g., [12]. However, it remains unclear how GNNs are actually encoding graph structure information into their vector representations, and whether there are theoretical advantages of GNNs compared to kernel based approaches.

## 1.1 Present Work

We offer a theoretical exploration of the relationship between GNNs and kernels that are based on the 1-WL algorithm. We show that GNNs cannot be more powerful than the 1-WL in terms of distinguishing non-isomorphic (sub-)graphs, e.g., the properties of subgraphs around each node. This result holds for a broad class of GNN architectures and all possible choices of parameters for them. On the positive side, we show that given the right parameter initialization GNNs have the same expressiveness as the 1-WL algorithm, completing the equivalence. Since the power of the 1-WL has been completely characterized, see, e.g., [1, 6], we can transfer these results to the case of GNNs, showing that both approaches have the same shortcomings.

Going further, we leverage these theoretical relationships to propose a generalization of GNNs, called  $k$ -GNNs, which are neural architectures based on the  $k$ -dimensional WL algorithm ( $k$ -WL), which are strictly more powerful than GNNs. The key insight in these higher-dimensional variants is that they perform message passing directly between subgraph structures, rather than individual nodes. This higher-order form of message passing can capture structural information that is not visible at the node-level.

Graph kernels based on the  $k$ -WL have been proposed in the past [8]. However, a key advantage of implementing higher-order message passing in GNNs—which we demonstrate here—is that we can design hierarchical variants of  $k$ -GNNs, which combine graph representations learned at different granularities in an end-to-end trainable framework. Concretely, in the presented hierarchical approach the initial messages in a  $k$ -GNN are based on the output of lower-dimensional  $k'$ -GNN (with  $k' < k$ ), which allows the model to effectively capture graph structures of varying granularity. Many real-world graphs inherit a hierarchical structure—e.g., in a social network we must model both the ego-networks around individual nodes, as well as the coarse-grained relationships between entire communities, see, e.g., [9]—and our experimental results demonstrate that these

hierarchical  $k$ -GNNs are able to consistently outperform traditional GNNs on a variety of graph classification and regression tasks. Across twelve graph regression tasks from the QM9 benchmark, we find that our hierarchical model reduces the mean absolute error by 54.45% on average. For graph classification, we find that our hierarchical models leads to slight performance gains.

## References

- [1] V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. “On the Power of Color Refinement”. In: *Symposium on Fundamentals of Computation Theory*. 2015, pp. 339–350.
- [2] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. “Neural Message Passing for Quantum Chemistry”. In: *ICML*. 2017.
- [3] M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Lecture Notes in Logic. Cambridge University Press, 2017.
- [4] W. L. Hamilton, R. Ying, and J. Leskovec. “Inductive Representation Learning on Large Graphs”. In: *NIPS*. 2017, pp. 1025–1035.
- [5] W. L. Hamilton, R. Ying, and J. Leskovec. “Representation Learning on Graphs: Methods and Applications”. In: *IEEE Data Engineering Bulletin* 40.3 (2017), pp. 52–74.
- [6] Sandra Kiefer, Pascal Schweitzer, and Erkal Selman. “Graphs identified by logics with counting”. In: *MFCS*. Springer. 2015, pp. 319–330.
- [7] T. N. Kipf and M. Welling. “Semi-Supervised Classification with Graph Convolutional Networks”. In: *ICLR*. 2017.
- [8] C. Morris, K. Kersting, and P. Mutzel. “Glocalized Weisfeiler-Lehman Kernels: Global-Local Feature Maps of Graphs”. In: *ICDM*. 2017, pp. 327–336.
- [9] M. E. J. Newman. “The structure and function of complex networks”. In: *SIAM review* 45.2 (2003), pp. 167–256.
- [10] N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. “Weisfeiler-Lehman Graph Kernels”. In: *JMLR* 12 (2011), pp. 2539–2561.
- [11] S. V. N. Vishwanathan, N. N. Schraudolph, R. Kondor, and K. M. Borgwardt. “Graph Kernels”. In: *JMLR* 11 (2010), pp. 1201–1242.
- [12] R Ying, J. You, C. Morris, X. Ren, W. L. Hamilton, and J. Leskovec. “Hierarchical Graph Representation Learning with Differentiable Pooling”. In: *NIPS*. 2018.





Subproject B1  
Analysis of Spectrometry Data with  
Restricted Resources

Jörg Ingo Baumbach

Jörg Rahnenführer

# Peak position alignment for two MCC-IMS devices

Salome Horsch

Statistical Methods in Genetics and Chemometrics

Technische Universität Dortmund

salome.horsch@tu-dortmund.de

The automated analysis of MCC-IMS raw measurements has become feasible due to the development of algorithms tailored for this special kind of data. Some problems arise however with regard to broad applicability of these algorithms, for example when data from different devices is used. We show that two different devices can lead to different results and how the measurement of a reference component mixture allows for improved peak detection.

Diagnosing diseases by analyzing breath gas has become a wide field of interest. Using MCC-IMS for breath gas analysis offers several advantages to using other technologies. In contrast to technologies like GC-MS, the breath composition can be analyzed on the spot due to fast processing and comparatively small devices that can be transported or put up in places with limited spacing, e.g. in an operation room. One drawback is however, that automated evaluation is not established yet and that technical details complicate the analysis. A raw measurement consists of signal intensity values in two dimensions, the retention time and the inverse reduced mobility. Together, they form a heat map where peak regions represent a certain compound in the measured air sample.

For the automated peak detection we use the algorithm combination that showed best subsequent classification results in [1], called SGLTR-DBSCAN. In this combination, SGLTR accomplishes the detection of single peaks in a raw measurement. Afterwards, DBSCAN compares the peak positions from all available raw measurements and decides, which peaks from several raw measurements are likely to be originating from the same metabolite in the corresponding breath gas. This decision is based on how close the peak positions from several measurements are and results in so-called consensus peaks. Single peaks are either assigned to consensus peaks or discarded when there are not enough single peaks close to it. If a raw measurement doesn't have a single peak that can be

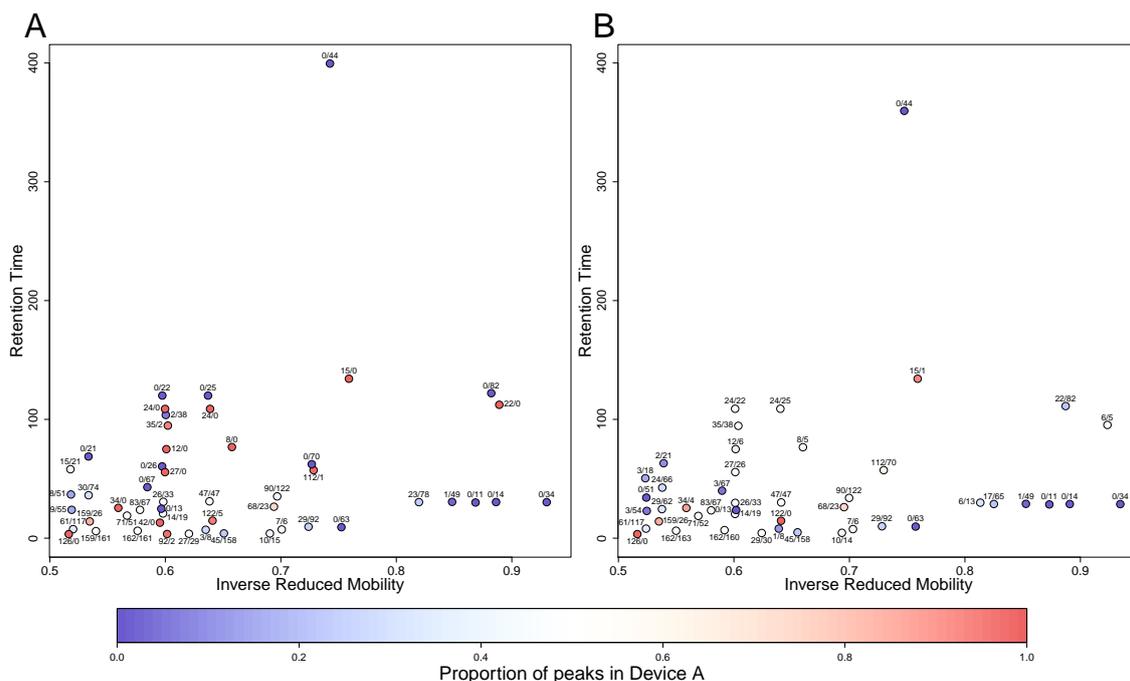


Figure 1: Number of single peaks from both devices corresponding to the displayed consensus peaks. Red dots show peaks appearing more often on device A, blue dots peaks appearing more often on device B. White colors mark peaks with approximately equal distributions. **A**: Automatic peak detection with SGLTR-DBSCAN. **B**: Automatic peak detection with inserted peak alignment after the SGLTR step.

assigned to a certain consensus peak, the value for the consensus peak regarding this raw measurement is recorded as zero. Finally, the dataset with raw measurements as observations and consensus peaks as features can be used for further analysis.

In order to assess differences between several devices, we use data from our study on healthy subjects. The breath of 49 probands was measured twice, once on each of the two devices A and B. For other purposes, every person drank a glass of orange juice before the second measurement. Accompanying each breath gas analysis, the room air was measured as well as a measurement that was undertaken during a cleaning flush. This results in 294 measurements corresponding to the breath samples. Additionally, cleaning was sometimes done in between study days, such that in total, 345 MCC-IMS measurements were taken. Of those, 172 were obtained from device A, the remaining 173 from device B.

The positions of the consensus peaks as determined by the automatic algorithm are shown in figure 1A. Since the automatic algorithm assigns the value zero for a raw measurement when no peak was detected at this position the presence of metabolites in measurements taken on the two devices can easily be evaluated by counting the number of raw measure-

ments having a value unequal to zero. Since both devices measure approximately similar conditions (observations were assigned to the devices by stratified randomization), the number of raw measurements containing a certain metabolite should be approximately equal. In figure 1A however, it can be seen that many peaks are found way more often on one of the two devices. Sometimes peaks of opposite proportions lie close to each other, suggesting that the same metabolic compound in the gas sample results in slightly shifted single peak positions on the two devices. As a result, the peaks form two separate clusters, i.e. consensus peaks in the clustering step.

To solve this problem, reference gas measurements (artificial gas probes with known composition) for the two devices were compared. The two corresponding datasets contained five matching components and their peak positions had already been determined manually. The software VisualNow allows setting values for stretching the axes but they have to be set manually or using exactly two "old" and two "new" peak positions for determination of the stretching factors. It is impossible however, to conveniently use more data for a more stable alignment in an automated fashion. For that reason, we instead suggest to use all components that were measured on both devices. Their peak positions can be aligned by applying a linear regression for each of the two dimensions (retention time and inverse reduced mobility), regressing the positions from one device on the positions from the other device, which thus serves as reference device.

Inserting this correction step in between the automatic single peak detection and the consensus peak clustering, the single peaks are shifted to similar locations before they are clustered. The results are shown in figure 1B. The resulting consensus peaks are much more plausible with respect to the numbers of corresponding single peaks from both devices. Especially in the area of inverse reduced mobilities between 0.60 and 0.65 many dots with rich colors (i.e. unequally distributed proportions on both devices) from figure 1A turn into fewer dots with brighter shades (i.e. equally distributed proportions on both devices) in figure 1B.

In addition to the peak position alignment we use data that has also been scaled with respect to the device. From now on, only breath data is used whereas room air and flush measurements are discarded (they were only used for stable peak detection). The automatic algorithm denotes peak intensities in raw measurements as zero, when no single peak was detected in the correspondent raw measurement at this position. Since zeros occur quite often in the resulting datasets, common normalization to mean zero and variance one separately for each device, results in counterintuitive results for the values that are zero. Since the sample means and variances are generally different for both devices, the zeros from both devices take different values after the normalization, although in both cases the metabolite was not detected and thus should have the same value after the standardization, too. Therefore, our scaling normalizes all non-zero values to mean zero and variance one for each device but the zero-values are shifted equally by subtracting the maximum of both sample means and dividing by the minimum sample

variance (thus guaranteeing, that zero-values remain smaller after the scaling than other values that were larger before).

The effects of both methods, peak position alignment and scaling, can also be observed when testing univariately for differences between means, using the device as the grouping variable. Since each person is measured once on each device, a paired test is used (here: Wilcoxon test). Table 1 shows the number of univariately significant metabolites in each scenario (with/without alignment and with/without scaling). In order to account for multiple testing, adjusted p-values according to the method of Bonferroni-Holm are reported (in parenthesis the numbers without p-value adjustment).

Table 1: Number of univariately significant metabolites for differences of mean for the two devices. A paired Wilcoxon test was used. The quantities are shown for data without peak position alignment as well as with peak position alignment each with or without scaling with respect to the device. The numbers in parenthesis concern unadjusted p-values whereas the others are based on adjusted p-values using the method of Bonferroni-Holm.

# significant metabolites	without scaling		with scaling	
	adjusted	unadjusted	adjusted	unadjusted
without alignment ( $p = 51$ )	24	(27)	11	(20)
with alignment ( $p = 46$ )	18	(23)	7	(14)

As can be seen in Table 1, the alignment step reduces the number of significant metabolites from 24 to 18 for unscaled and from 11 to 7 for scaled data. This coincides with the reduced number of found consensus peaks that decreases from 51 to 46 when the alignment is performed. Additionally, scaling the data with regard to the device, differences are less frequently observed. This also shows that this scaling doesn't remove all differences between the devices since the procedure of treating zeros differently maintains differences in the occurrence of a metabolite between the devices. These differences could result from different materials that are built into the devices.

## References

- [1] S. Horsch, D. Kopczynski, E. Kuthe, J. I. Baumbach, S. Rahmann, and J. Rahnenführer. A detailed comparison of analysis processes for mcc-ims data in disease classification - automated methods can replace manual peak annotations. *PLoS One*, 12, 2017.





## Subproject B2

Resource optimizing real time analysis of  
artifactual image sequences for the  
detection of nano objects

Roland Hergenröder

Frank Weichert

# Cache-Aware Optimization for Decision Tree through Tree Framing

Kuan-Hsun Chen

Lehrstuhl für Technische Informatik und Eingebettete Systeme  
Technische Universität Dortmund  
kuan-hsun.chen@tu-dortmund.de

Nowadays, the ongoing integration of machine learning models into everyday life also demand the evaluation to be extremely fast and in real-time. The implementation of the model application must take the characteristics of the executing platform into account. Due to "memory wall" phenomenon, hierarchical cache memories have been widely used in most modern architectures to hide the memory latencies. For such embedded systems, the adoption of cache memories becomes really critical. This report briefly go through the idea how can we optimize the execution of Decision Trees through tree framing approaches.

## 1 Introduction

Decision Trees (DT) form the basis of many ensemble methods, such as Random Forests (RF) or Extremely Randomized Trees (ET). As one of the best black-box methods available offering high accuracy with only a few parameters to tune [3,4], RF are widely used in various research domains, e.g., Physicist uses Random Forest to filter measurements of monitoring the sky for gamma rays. For these methods to work best, trees should be as large as possible.

In the last year lately, we started to investigate the data locality issues in most of the implementation of DT. We observed that, the memory excess pattern of DT (even in its basic unit), has no data locality at all. Especially on embedded systems, which are used to perform real-time calculations, the available resource is usually limited. Such a case, i.e., big trees but small caches, really challenges the data and the instruction

cache of CPUs and thus demand a more careful memory layout. Particularly for runtime considerations, it has been stated that caching behaviour determines the performance of implemented algorithms even more than algorithmic differences [5]. This motivates us to revisit the basis of computer architecture and the concept of cache design.

Finally, we come out several principled approaches and provide a comprehensive code generator adapting to particular parameters of the computer's memory and produce optimized code segments in [1]. In short, we consider a probabilistic view of DT model execution in CPUs, and introduce a theoretically well-founded memory layout. Throughout the evaluations, we can observe that our implementation offers a speed-up factor from 2 to 4 on average without changing the prediction accuracy of the model.

## 2 Ideas and Results

Following the probabilistic view of DT execution in [2], we model each comparison at node  $i$  as a Bernoulli experiment in which we will take the path towards the left child with probability  $p(i \rightarrow l(i))$  and respectively for the right child with  $p(i \rightarrow r(i))$ . It holds that  $p(i \rightarrow l(i)) = 1 - p(i \rightarrow r(i))$ . Based on this probabilistic arguments, we can further model the probability of each path during training by counting the number of samples at each node  $i$  taking the left and right path. We further assume the quality of the learned model to be satisfactory and do not change it. More details please refer to [2].

We specially focus on two typical implementations of DT: Native Tree (shown in 1), which uses a loop to iterate over each node of a tree within a continuous data structure, e.g., arranged by a one-dimensional array. and If-else Tree (shown in 2), which uses the split values of a tree are all hard-coded as constant values into the instructions.

```

struct Node {
    bool isLeaf;
    unsigned int prediction; // Predicted Label
    unsigned char feature; // Targeted feature
    float split; // Threshold
    unsigned short leftChild;
    unsigned short rightChild;
};
Node tree [] = {{0,0,0,8191,1,2},{0,0,1,2048,3,4},...}
bool predict(short const x[3]){
    unsigned int i = 0;
    while(!tree[i].isLeaf) {
        if (x[tree[i].f] <= tree[i].split) {
            i = tree[i].left;
        } else {
            i = tree[i].right;
        }
    }
    return tree[i].prediction;
}

```

Listing 1: Native tree structure in C++

```

bool predict(short const x[3]){
    if(x[0] <= 8191){
        if(x[1] <= 2048){
            return true;
        } else {
            return false;
        }
    } else {
        if(x[2] <= 512){
            return true;
        } else {
            return false;
        }
    }
}

```

Listing 2: Example for If-else structure in C++.

Although the if-else tree already unrolls the comparisons of a DT into conditional statements without indirect memory accesses, cache misses may still occur. To reducing compulsory cache misses, we propose to traverse all its paths and swap the children of every node  $i$  when  $p(i \rightarrow l(i)) \geq p(i \rightarrow r(i))$ . To reduce capacity and conflict cache misses, we define a computation kernel containing those nodes which are used most of time, and leverage goto to break the tie of branches.

For the native-tree, we propose to reduce the amount of memory each node needs as much as we can to reduce compulsory cache misses, i.e., abandon the `isLeaf` and `prediction` field of the native solution, but store the prediction of the left (right) child directly in the respective fields `left` (`right`) if it is a leaf node. Furthermore, we propose an approach to arrange the location of nodes stored in memory for creating the data locality by using the probabilistic information.

To deploy extensive evaluations, we have performed 1800 different experiments by training Decision Trees (DT), Random Forests (RF), and Extremely Randomized Trees (ET) on 12 different data-sets with varying tree-depths to generate the aforementioned implementations for different architectures, i.e., X86, PPC and ARM CPUs. The results on X86 show that if-else trees should be adopted and the improvement can be up to 3 times compared to the naive native tree. For ARM, we notice that there is no clear winner of implementations for larger trees. Interestingly, the optimized native tree implementation can even outperform the standard if-else tree due to the large size of D-Cache.

### 3 Outlook

When machine learning meet embedded systems, the efficiency of black-box methods is required to be reconsidered under the assumption of limited resources regarding to time and memory spaces. In this report, I briefly go through the idea how we advocate

for optimizing the evaluation of learned models in the real-time setting in [1]. Since tree ensembles in general and specifically Random Forests are among the most efficient black-box methods available, we looked at the application of decision trees from a computer-architectures point of view.

Our findings of this work show that implementing RF can be challenging and need to be done carefully, since the hierarchical cache memories in the underlying computer architectures can be really important especially for embedded systems. The propose approaches only exploit some features regarding memory hierarchy. Implementations that consider other architectural features like scratchpad, out-of-order execution, hyper-threading, branch prediction, etc. may further improve the performance, which are considered as future work. Recently we start to investigate the usage of random forests under the consideration of concept drifts. We notice that indeed the implementation of well-known approaches is worth to be revisited and redesigned.

## References

- [1] S. Buschjäger, K. Chen, J. Chen, and K. Morik. Realization of random forest for real-time evaluation through tree framing. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 19–28, Nov 2018.
- [2] S. Buschjäger and K. Morik. Decision tree and random forest implementations for fast filtering of sensor data. *IEEE Transactions on Circuits and Systems I: Regular Papers*, PP(99):1–14, 2017.
- [3] R. Caruana, N. Karampatziakis, and A. Yessenalina. An empirical evaluation of supervised learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 96–103. ACM, 2008.
- [4] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim. Do we need hundreds of classifiers to solve real world classification problems? *J. Mach. Learn. Res*, 15(1):3133–3181, 2014.
- [5] Siegfried Nijssen and Joost N. Kok. Frequent subgraph miners: Runtime don't say everything. In *Proceedings of the International Workshop on Mining and Learning with Graphs (MLG 2006)*, pages 173–180, 2006.





Subproject B3  
Data Mining on Sensor Data of Automated  
Processes

Jochen Deuse      Katharina Morik  
Petra Wiederkehr

# Learning Ensembles in the Presence of Imbalanced Classes.

Amal Saadallah

Lehrstuhl für Künstliche Intelligenz, LS 8

Technische Universität Dortmund

amal.saadallah@tu-dortmund.de

In many use cases in the subproject B3, rare events and unusual behaviors, such as process failure or instability in machine engineering, are usually represented with imbalanced data observations. In other words, one or more classes, usually the ones that represent such events, are underrepresented in the data set. This issue, known to the Data Mining community as the class imbalance problem. However, class imbalance is known to induce a learning bias towards majority classes which implies a poor detection of minority classes. Thus, we propose a new probabilistic ensemble method to handle class imbalance explicitly at training time.

## 1 Introduction

Several machine learning approaches have been proposed over the past decades to handle the class imbalance problem. Most of which are based on resampling techniques, cost sensitive learning, and ensemble methods [1, 2, 4]. In this paper, we address the problem of class imbalance via an ensemble method. Our method falls within resampling-based ensembles, as before constructing the ensemble, it uses random sampling for designing scenarios of class imbalance on unseen test data and for generating balanced training sets to train single classifiers by optimizing a well-defined loss function for each scenario. Our method combines (estimated) knowledge about the class imbalance in the data distribution with random sampling to handle the class imbalance problem.

## 2 A Mixture of Experts for Class Rebalancing

Usually, when we use random sampling for balancing the training set  $\mathcal{D}$ , we evaluate the quality of sampled instances and their ability in solving the classification problem on an independent validation/calibration set  $\mathcal{D}_{\text{Val}}$ . Whenever we come to the deployment part of the model, we have obviously a new unseen test set. Even based on the strong assumption that data points are coming from the same distribution, the class imbalance scenario (i.e. class imbalance ratio  $IR$ : the ratio between the number of majority and minority class samples) on this new set can be different from the one on the training set and the calibration set, based on which the sampling procedure was optimized.

When comparing the possible class imbalance ratio of the test set to the one of the training set, three possible scenarios are then present (i.e. either we have similar  $IR$ , either higher or lower). To prepare our classifier for these situations, we construct three new data sets  $\mathcal{D}_{\text{Val}}^>$ ,  $\mathcal{D}_{\text{Val}}^{\approx}$ ,  $\mathcal{D}_{\text{Val}}^<$ . In the first case, the class  $y = 1$  (i.e. we assume its the label of the majority class) is actually more likely than the data suggests, we should create a subsample  $\mathcal{D}_{\text{Val}}^>$  from  $\mathcal{D}_{\text{Val}}$ , such that  $\tilde{\mathbb{P}}^>(\mathbf{Y} = 1) = \tilde{\mathbb{P}}(\mathbf{Y} = 1) + \alpha$ , where  $\tilde{\mathbb{P}}^>$  is the class distribution in  $\mathcal{D}_{\text{Val}}^>$  and  $\alpha$  is a positive probability offset—we refer to this process as *rebalancing*. The third case is symmetrical: we have  $\alpha < 0$  and we shall subsample  $\mathcal{D}_{\text{Val}}^<$  from  $\mathcal{D}_{\text{Val}}$ . Finally, for the case (2), we subsample the set  $\mathcal{D}_{\text{Val}}^{\approx}$  via stratified proportionate allocation—the class distribution will approximately match the class distribution in  $\mathcal{D}_{\text{Val}}$  and  $\alpha \approx 0$ .

We subsample the validation set  $\mathcal{D}_{\text{Val}}$  to generate three specific sets  $\mathcal{D}_{\text{Val}}^>$ ,  $\mathcal{D}_{\text{Val}}^{\approx}$ ,  $\mathcal{D}_{\text{Val}}^<$ . For each case  $\mu \in \{>, \approx, <\}$ , we learn a model  $m^\mu$  on the training data and use the data in the validation set  $\mathcal{D}_{\text{Val}}^\mu$  to refine the learning process. This procedure yields three families of models, namely  $M^>$ ,  $M^{\approx}$  and  $M^<$ .  $M^\mu$  may contain one or many models  $m^\mu$  depending on the chosen optimization procedure e.g., repeating the subsampling process multiple times or optimizing different quality metrics  $q$ . The resulting models will later be combined to form an ensemble, which is prepared to handle different scenarios of class discrepancy and different quality metrics  $q$ . The ensemble weights can be considered as unobserved random variables and apply expectation-maximization to estimate them [3]. However, such techniques assume the existence of a context and require the classifier accuracy to be Lipschitz continuous. In contrast, we propose to interpret  $\mu \in \{>, \approx, <\}$  and  $q$  as random variables. Let  $m(\mathbf{x}) = \hat{y}$  be the prediction of a single model, the outcome of the ensemble is:

$$\mathbb{E}[m(\mathbf{x})] = \sum_{\mu \in \{>, \approx, <\}} \mathbb{P}(\mu) \sum_{q \in \mathcal{Q}} \mathbb{P}(q | \mu) \sum_{t=1}^T \mathbb{P}(t | \mu, q) m_{q,t}^\mu(\mathbf{x})$$

where  $\mathcal{Q}$  is a set of quality metrics, e.g.,  $\mathcal{Q} = \{\text{precision, recall}\}$ , and  $T$  is the number of models that we have generated for each combination of  $\mu$  and  $q$ —our experimental results will show that small values of  $T$  suffice. The values of  $\mathbb{P}(\mu)$ ,  $\mathbb{P}(q | \mu)$ , and  $\mathbb{P}(t | \mu, q)$  control how the specific models  $m_{q,t}^\mu$  influence the final outcome. The probability  $\mathbb{P}(\mu)$

expresses our knowledge on how the class distribution in the training set differs from the true one. The probabilities  $\mathbb{P}(q \mid \mu)$  can be interpreted as the *importance of the quality measure*  $q$  for the actual prediction task. As such, the importance depends on  $\mathcal{D}$ ,  $\mathcal{Q}$ , or even on  $\mu$ , and cannot be derived in general. Finally,  $\mathbb{P}(t \mid \mu, q)$  are the classical (normalized) ensemble weights. Since each weak model depends on  $\mu$ ,  $q$ , and  $t$ , the outcome may be interpreted as a mixture of mixtures of experts.

### 3 Class Imbalance and the Probability Offset

Let us now put some light on how to choose the *probability offset*  $\alpha$  for the *rebalancing* process. At a first glance, choosing the probability offset  $\alpha$  might seem infeasible. How should one even guess which  $\alpha \in (0; 1)$  is appropriate? Surprisingly, it turns out that class imbalance *simplifies* this problem! In fact, the stronger the class imbalance, the tighter is the range for reasonable probability offsets. Let us formalize this result.

**Lemma 1 (Probability of Probability Offsets)** *Let  $\mathcal{D}$  be a data set with empirical class distribution  $\tilde{\mathbb{P}}(y)$ . We denote the class ratio w.r.t. the minority class by  $r = \min \left\{ \frac{\mathbb{P}(\mathbf{Y}=0)}{\mathbb{P}(\mathbf{Y}=1)}, \frac{\mathbb{P}(\mathbf{Y}=1)}{\mathbb{P}(\mathbf{Y}=0)} \right\}$ . There exists  $c, c'$  such that for all  $\alpha > 0$ , the probability of the event*

$$|\tilde{\mathbb{P}}(\mathbf{Y} = 0) - \mathbb{P}(\mathbf{Y} = 0)| \geq \alpha$$

*is upper bounded by*

$$\exp \left( -\frac{\alpha^2 c}{r + \alpha c'} + \ln 2 \right).$$

The lemma tells us that the probability that  $\alpha$  is large decreases exponentially fast as a function  $\alpha$ , damped class imbalance which is measured by  $r$ . As an example, assume that our data set contains  $N = 1000$  samples and the classes are perfectly balanced, i.e.,  $\mathbb{P}(\mathbf{Y} = 1) = 1/2$ . In this case, the lemma asserts that the probability that  $\tilde{\mathbb{P}}(\mathbf{Y} = 0)$  and  $\mathbb{P}(\mathbf{Y} = 0)$  deviate by at least 5% is upper bounded by 58.487%. Notice that the randomization happens over all data sets—this means that for more than half of all data sets, the estimated and true class frequency will deviate by more than 5%. Now, let us consider class imbalance. We assume  $\mathbb{P}(\mathbf{Y} = 1) = 0.9$ , i.e., 10% of all data points belong to the minority class. By applying Lemma 1 to this scenario while keeping all other quantities fixed, we see that a deviation of more than 5% happens in at most 0.011285% of all data sets. A deviation of at least 2.5% occurs for at most 14.615% of all data sets. Thus, we see that large  $\alpha$  values (say,  $> 5\%$ ) are very unlikely in the presence of class imbalance. Our approach does not try to fix a specific flaw of imbalanced data sets. Instead, we try to fix a flaw that is inherent in *all* classification tasks—however, fixing this flaw on imbalanced data sets is much easier, due to the heavily limited range of possible

probability offsets  $\alpha$ . Our experiments on six data sets underpin these theoretical insights and show that our framework outperform many State-of-the-Art methods in handling imbalanced data sets.

## 4 Future Work

As a future work, we plan to test our framework on highly imbalanced data sets using different accuracy measures. In addition, we plan to enhance the accuracy of our ensemble by promoting diversity among its members. This will be achieved by trying to select different base-line classifiers for each component of the ensemble in an informed manner and to combine it together with the adequate random sampling strategy.

## References

- [1] Mikel Galar, Alberto Fernandez, Edurne Barrenechea, Humberto Bustince, and Francisco Herrera. A review on ensembles for the class imbalance problem: bagging-, boosting-, and hybrid-based approaches. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(4):463–484, 2012.
- [2] Guo Haixiang, Li Yijing, Jennifer Shang, Gu Mingyun, Huang Yuanyue, and Gong Bing. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications*, 73:220–239, 2017.
- [3] Erfan Soltanmohammadi, Mort Naraghi-Pour, and Mihaela van der Schaar. Context-based unsupervised data fusion for decision making. In *Proc. of the 32nd ICML*, pages 2076–2084, 2015.
- [4] Yanmin Sun, Andrew KC Wong, and Mohamed S Kamel. Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence*, 23(04):687–719, 2009.

# Quality inspection in manufacturing by predictive machine learning models

Jacqueline Schmitt  
Institute of Production Systems  
Technische Universität Dortmund  
jacqueline.schmitt@ips.tu-dortmund.de

Due to growing competitive pressure, the quality of products is becoming an increasingly important success factor for manufacturing companies. In order to be able to guarantee high quality, extensive inspections are indispensable which are connected to high resource consumption. In this context, the application of machine learning (ML) techniques for predicting quality-related product features is becoming increasingly important. This technical report describes the requirements for classical measurement and inspection processes and how these must be transferred to predictive ML-models.

## 1 Introduction

Quality has become a key competitive factor for manufacturing companies. In order to ensure high-quality products, extensive inspections are unavoidable. Due to the high resource consumption of those non-value-adding processes, the development of alternative approaches is of great interest in research and industrial application. In the context of increasing digitalization, the amount of data available in manufacturing processes has grown rapidly in recent years. This enables data-driven approaches such as the application of machine learning techniques in industry. One field of interest thereby is the field of predictive quality where ML models are used to predict the expected final product quality based on previously recorded process parameters. Initial studies have shown promising results in various industries such that the gap between research and industry must now be closed in order to integrate ML models into the daily toolbox of manufacturing companies. In this context, the suitability of ML models as measuring equipment must be proven in order to replace classical quality inspection processes in the mid-term future.

## 2 Measuring system suitability

A basic prerequisite for measuring quality is a capable measuring system or measuring equipment and a reliable measuring method that allows the quality characteristic to be recorded correctly. Consequently, measuring equipment must meet certain requirements, their suitability must be demonstrated and the capability of the measuring process must be analysed.

### 2.1 Measuring equipment requirements

The variation of the measurement results (total variation) is on the one hand due to the variation of the measurement process, which as an independent process is afflicted with a certain scatter, and on the other hand on the variation of the measuring system. The basic objective is to keep the measurement errors and deviations resulting from the variation of the measurement system to a minimum. Just like the actual measured values, the measurement deviations are influenced by a large number of factors and can lead to incorrect conclusions when considering the overall variation. Consequently, measuring equipment and systems have to meet certain requirements in order to provide reliable measurement results: sufficient resolution, accuracy and precision, repetition and comparison precision, stability, linearity [4, 5].

### 2.2 Measurement system analysis

The measuring system analysis serves the analysis of the capability of measuring equipment and measuring systems. A new piece of inspection equipment is first subjected to an incoming inspection in which the equipment is calibrated with a standard in the measuring room [3]. The analysis of the equipment's capability then takes place at the place of use under the measurement conditions given there. Provided that the measured values are normally distributed, various methods can be applied using standards or production parts which are measured repeatedly [3]:

**Method 1:** Verification of the capability of the measuring equipment

**Method 2:** Proof of the capability of a measurement process with operator influence (Gauge R&R Study)

**Method 3:** Proof of the ability of a measurement process without operator influence (R&R Study)

**Method 4:** Verification of the measurement stability of a measurement process

**Method 5:** Proof of capability for inspection processes for attributive characteristics

## 2.3 Process capability

In addition to the capability of the measuring system, process capability is also an important prerequisite for measuring the quality of a product. Process capability is thereby the basic ability of a process to meet customer requirements. A distinction can be made between 3 types of process capability [1–3]:

1. machine capability ( $C_m, C_{mk}$ )
2. short-term process capability ( $C_p, C_{pk}$ )
3. long-term process capability ( $P_p, P_{pk}$ )

The calculation of the indices is analogous for all 3 types and depends on the data used. If the data is collected under ideal conditions, only the machine capability can be proven. If, on the other hand, the data is collected under real production conditions, short-term or long-term process capability can be demonstrated depending on the length of the observation period.

## 3 Transfer on predictive ML-models

In order to replace classical measurement and testing processes with predictive ML-models in the mid-term future, their equivalence must first be proven. For this purpose, the requirements placed on classical measuring and inspection equipment must be transferred to the models and their fulfilment must be proven.

### 3.1 General idea of predictive modeling

Predictive modeling is an approach that uses mathematical and computational methods to predict an event or outcome based on historic data. A ML-model describes the relationship between input variables and model output and allows to predict an outcome at some future state or time based upon changes in the model input. In industrial applications, these models enable the prediction of quality-related product characteristics on the basis of previously recorded process parameters. This allows to relieve resource-intensive inspection processes with constant reliability of high-quality products.

### 3.2 Model suitability for quality inspection

The requirements for classical measuring systems also apply for the quality observation by means of predictive ML-models. However, the requirements cannot be transferred directly but their purpose has to be understood and converted. The suitability of models highly depends on the quality of the input data which therefore has to be specifically investigated. Additionally, the underlying algorithm as well as the model structure and parameters influence the applicability and must therefore be taken into account for every individual application.

The decision on which method of measurement system analysis to choose depends on the mode of integration of the model. If the judgement on quality is purely model-based, the operator influence can be neglected. Depending on the feature type of the label, continuous or discrete, regression or classification models will be used and accordingly the measurement system analysis method for attributive characteristics has to be applied in the latter case.

The process capability analysis can be performed analogously to the classical application whereas the time window of the recorded data determines the type of process capability. While the machine capability represents the model application without any noise factors, the short-term and long-term process capability include noise factors that may influence the model performance negatively.

## 4 Research outlook

Further research is going to investigate the application of ML-based predictive models as measurement equipment for quality inspection in automated manufacturing processes. For this purpose, equivalent requirements and verification methods must be developed in order to increase confidence in such test equipment. Furthermore, the inclusion of mode-based inspection equipment leads to significant changes in inspection planning and inspection equipment management. Especially the hybrid application of classical measurement and inspection equipment and model-based inspection poses great challenges that have to be addressed in further investigations.

## References

- [1] Claudia Brückner. *Qualitätsmanagement: Das Praxishandbuch für die Automobilindustrie*. Hanser Verlag, München, 2011.
- [2] Edgar Dietrich. Eignungsnachweise für messprozesse. In Tilo Pfeifer and Robert Schmitt, editors, *Masing Handbuch Qualitätsmanagement*, pages 665–684. Carl Hanser Fachbuchverlag, s.l., 2014.
- [3] Claus P. Keferstein, Michael Marxer, and Carlo Bach. *Fertigungsmesstechnik: Alles zu Messunsicherheit, konventioneller Messtechnik und Multisensorik*. Springer Vieweg, Wiesbaden, 9., überarbeitete und erweiterte auflage edition, 2018.
- [4] Tilo Pfeifer and Robert Schmitt. *Fertigungsmesstechnik*. Oldenbourg, München, 3., überarb. und erw. aufl. edition, 2010.
- [5] Robert Schmitt and Tilo Pfeifer. *Qualitätsmanagement: Strategien - Methoden - Techniken*. Hanser, München, 5., aktualisierte auflage edition, 2015.

# **Collaborative Research Center SFB 876 - Adaptive soft sensors - Dealing with concept drift in sensor data**

Mario Wiegand  
Institut für Produktionssysteme  
Technische Universität Dortmund  
mario.wiegand@ips.tu-dortmund.de

Production processes constantly change their state and characteristics, so that recorded process and product data underlie concept drift. To deal with concept drift several approaches exist in the literature. In this report we present an overview on these approaches and identify need for further research. In the third phase of the Collaborative Research Center 876 we want to develop new methods that can detect concept drift and update the models accordingly.

## **1 Introduction**

Within the Collaborative Research Center 876 project B3 deals with the time-constrained analysis of sensor data in production processes using machine learning techniques. In the first two phases a quality prediction system could be developed allowing for an in process-prediction of the product's quality features based on real-time data from the production process. The advantage of this approach is, that features, which are difficult to measure online, can be monitored continuously. As a result, quality problems can be detected early in the process directly after their emergence, so that actions can be taken immediately, in order to improve process and product quality. Models, which allow for the prediction of difficult-to-measure variables by easy-to-measure variables in production processes are called soft sensors or virtual sensors.

In the first two phases of the project the models were trained offline on the basis of batch data and retained stationary. However, when using soft sensors in real-world processes, it is of the utmost importance to guarantee the long-term operability and quality of the sensors. In this context, a particular challenge is the process dynamics, which leads to process conditions and process states constantly changing over time. Causes for these time-dependent changes are for instance tool wear, changes in input materials, products with varying quality requirements and changing environmental conditions (e.g. temperatures, humidity). In the machine learning literature such state changes are summarized under the term concept drift.

In the third phase of the project our aim is to develop advanced incremental learning methods that are able to adapt to the evolution of sensor data over time with online efficient memory length choice. In the soft sensor literature a variety of different adaption approaches to handle concept drift have been proposed. Thus, in the first step, we did a literature search, in order to get an overview on existing approaches and identify gaps for further research.

## **2 Existing approaches for model adaption**

The most common approaches for model adaption in the literature are the moving window approach, recursive adaption and just-in-time learning:

The moving window approach aims to update the model on the basis of a set of selected data points. The set is selected to represent the current concept of the process as precisely as possible. It is assumed that the most recent observations best reflect the current concept, so that they form the subset [3]. In the majority of the papers, a sliding window is moved along the data, which includes the newest samples and simultaneously removes the oldest samples from the dataset [7]. In this way, the model is continuously updated based on the most recent data. The moving window approaches are either based on an incremental model adaption, in which the model is updated for each newly recorded sample, or on a blockwise adaption, where the model update is done for a batch of samples together [3]. The adaption interval, i.e. the number of samples recently included in the time window before the model is re-trained, is called the step size [5]. Using incremental model adaption, the step size is one. A high step size leads to a lower update frequency of the model and thus to a lower computational effort, but there is a danger that the model is adapted too late, so that process conditions have already changed and predictions become imprecise. A disadvantage of the moving window approach is, that it has difficulties in adapting the model to abrupt process changes, since the model is always influenced by older samples in the set of selected data points. This applies in particular to large time windows and step sizes. Furthermore, the required storage capacity of moving window

approaches can become a problem, especially in resource-constrained applications, since all data of the time window must be kept in memory [6].

In contrast to the moving window approach described above, recursive adaptation methods do not update the model by completely re-training the model, but by recursively adapting the existing model using the newly recorded samples [4], e.g. by updating the covariance matrix. Most of the papers describe recursive formulations of principal component analysis or partial least squares regression. However, there are also papers dealing with a recursive support vector machine or a recursive neural network. Similar to the moving window approach, recursive adaptation is disadvantageous in handling abrupt process changes as the model is always influenced by older samples of the dataset.

The approach of just-in-time learning (JITL) is based on ideas from database technology and local modelling [1]. The general idea of JITL is similar to the moving window approach because the model is adapted by updating the training dataset. In contrast to the moving window approach, the most recent samples are not used for modelling, but rather the samples with the greatest similarity to the new samples [6]. The prerequisite for this is the creation and maintenance of an extensive database in which all historical and new samples are stored. As soon as a new sample arrives for which the label is to be predicted, relevant samples from the database with a high similarity are determined and used for modelling. The model is then used for predicting the label and is discarded subsequently. So in JITL a new local model is trained online on demand, which is only valid for the current sample. JITL approaches offer the possibility to adequately model time-varying and non-linear processes through the online training of local models. Both continuous and abrupt process changes can be modelled appropriately [2]. A disadvantage, however, is the high computational effort associated with the similarity search and the frequent model training. In addition, the historical database must completely contain all possible process states in order to be able to guarantee a high prediction quality of the local models in the case of process changes.

### **3 Direction of research in the third phase**

The literature review shows that a lot of work has been done in the context of adaptive soft sensors. However, the existing approaches still have problems and disadvantages that have to be overcome. In addition, most of the approaches described in the literature are based on an automatic update of the soft sensor model. If the update frequency is low, this may lead to a late reaction to process changes, so that the model is temporarily invalid. On the other hand, a high update frequency or even a continuous model adaptation leads to a very high computational effort, which usually cannot be managed in real time. In order to adapt the model only on demand, when a process change occurs, we want to develop an approach that is able to detect process changes immediately, so that the

model can be updated without time delay. We will focus on building models that are able to detect changes, distinguish between noise and drifts and are adaptive to changes, while remaining robust to noise.

## References

- [1] Zhiqiang Ge and Zhihuan Song. Online monitoring of nonlinear multiple mode processes based on adaptive local model approach. *Control Engineering Practice*, 16:1427–1437, 2008.
- [2] Huaiping Jin, Chen Xiangguang, Jianwen Yang, Li Wang, and Lei Wu. Online local learning based adaptive soft sensor and its application to an industrial fed-batch chlortetracycline fermentation process. *Chemometrics and Intelligent Laboratory Systems*, 143:58–78, 2015.
- [3] Petr Kadlec, Ratko Grbic, and Bogdan Gabrys. Review of adaptation mechanisms for data-driven soft sensors. *Computers and Chemical Engineering*, 35:1–24, 2011.
- [4] Manabu Kano and Morimasa Ogawa. The state of the art in chemical process control in japan: Good practice and questionnaire survey. *Journal of Process Control*, 20:969–982, 2010.
- [5] Bo Lu, Ivan Castillo, Leo Chiang, and Thomas F. Edgar. Industrial pls model variable selection using moving window variable importance in projection. *Chemometrics and Intelligent Laboratory Systems*, 135:90–109, 2014.
- [6] Agus Saptoro. State of the art in the development of adaptive soft sensors based on just-in-time models. *Procedia Chemistry*, 9:226–234, 2014.
- [7] Le Yao and Zhiqiang Ge. Moving window adaptive soft sensor for state shifting process based on weighted supervised latent factor analysis. *Control Engineering Practice*, 61:72–80, 2017.





Subproject B4  
Analysis and Communication for Dynamic  
Traffic Prognosis

Thomas Liebig      Michael Schreckenberg  
Christian Wietfeld

# Experimental Evaluation System for 5G New Radio mmWave Communications in UAV Application Scenarios

Karsten Heimann

Lehrstuhl für Kommunikationsnetze  
Technische Universität Dortmund  
karsten.heimann@tu-dortmund.de

In the fifth generation of mobile communication, the demand for enhanced mobile broadband services with strong throughput requirements substantiates the utilization of new spectrum in the millimeter wave domain, where high bandwidths are available. With these novel frequencies like 28 GHz, highly directional antennas are expected to counteract the more delicate radio conditions. Phased arrays are particularly used to form a high gain pencil beam. A precise alignment of transmitter and receiver is required to profit from the essential antenna gains. This is supposed to be a challenging task especially for network participants with a high mobility like Unmanned Aerial Vehicles (UAVs). For this reason, an experimental evaluation of the pencil beam alignment is carried out in a UAV context. Results show, that an alignment within a given error margin allows for a stable air-to-ground connection.

## 1 Motivation

Complementary machine-type application scenarios are to be addressed by the fifth generation of mobile communication (5G) beside the conventional human-centric ones. Especially the growth of so called enhanced Mobile Broadband (eMBB) services longs for new spectrum resources to enable higher data rates. Though the millimeter wave (mmWave) domain offers the required bandwidth, but the drawback of this higher frequencies comes with the more challenging radio environment. Not only the increasing path loss, but also

the reduction of the indoor penetration capability and the demand for line-of-sight (LOS) constitute harsh conditions. In order to cope with this, high gain antennas are believed to more than compensate the drop of the link budget. Particularly beamforming concepts appear to be likely suited to grant high gains by maintaining a considerable flexibility.

Considering eMBB use cases, the context of Unmanned Aerial Vehicles (UAVs) may even put a new spin on applications like high-resolution video surveillance or ad-hoc mobile network coverage. Since conventional mobile networks are designed for a comprehensive ground coverage, UAVs currently suffer from poor coverage at altitudes. The disadvantageous effect of base stations' downtilt is illustrated in [1]. In contrast to that, upcoming mmWave base stations are expected to utilize phased array antennas and as a consequence thereof to achieve high beamforming or antenna gains. A precise alignment of a penciled main lobe ("pencil beam") is therewith viable towards any direction, so that an on demand coverage for arbitrary network subscribers including UAVs may be reached. Figure 1 depicts the proposed beam steering capabilities of a ground station and a UAV. While the ground station may be equipped with a phased array antenna to adjust the beam's direction by beamforming, the UAV may merely use a lightweight horn antenna which can be coarsely aligned mechanically due to its larger beam width.



Figure 1: The ground station applies beamforming, whereas the UAV needs to only coarsely align by azimuthal rotation.

## 2 Experimental Proof-of-Concept Evaluation

In our work [2], we focus on the beam alignment capabilities by experimentally evaluate the mmWave communication link between a UAV in flight and a ground station. To this end, the mmWave ground station is equipped with a 64 element phased array antenna, while a horn antenna mounted on the UAV is tethered to the mmWave R&D system. In Figure 2, the experimental setup is shown. It should be noted that the horn antenna has a larger half-power beam width of about  $54^\circ$  compared to about  $13^\circ$  of the phased array. Against this background, measurement results have shown, that a mechanical horn antenna alignment by azimuthal rotation of the UAV is feasible and essential during flight. In the course of this, the larger beam width relaxes the required alignment precision of the UAV's antenna. Nevertheless



Figure 2: Indicated pencil beam from ground station and horn antenna on UAV in flight

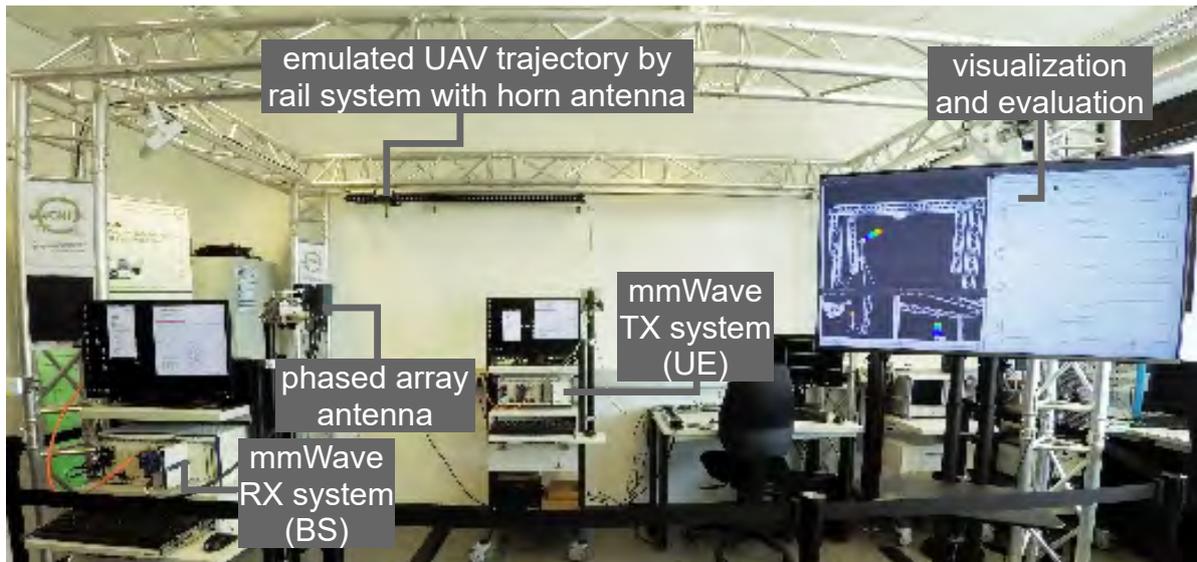


Figure 3: Laboratory test environment with rail system on truss for highly reproducible emulation of UAV trajectories.

a precise tracking of the pencil beam from the ground station towards the UAV needs to be realized within a misalignment margin of  $|\chi| < 5^\circ$  and a high accuracy with low angular noise is imperative (c.f. [2]).

### 3 The proposed 5G mmWave Antenna Tracking System

After the conceptual evaluation, a sophisticated laboratory testing and measurement environment needs to be designed to allow for highly reproducible and well-defined measurement procedures. For this reason, we decide to emulate the UAV's trajectory by a rail system as depicted in Figure 3:

A tooth belt stepper motor enables arbitrary acceleration patterns within a rail length of three meters, while a worm gear stepper motor synchronously ensures the proper alignment of the mounted horn antenna by means of a predefined base station position. Again, the horn antenna is tethered to the static mmWave system placed in the middle below the rail to promote preferably short HF cabling. On the left, the base station is deployed on a laboratory cart together with the phased array antenna. As a result of this, the base station antenna's physical pose can be adjusted to satisfy the application scenario to be investigated. On the right, a visualization and real-time evaluation entity displays the currently applied directional radio pattern of the phased array antenna as well as the pose of the emulated UAV, while the truss is implied as a reference (Figure 4). Additionally, the statistics about the UAV's pose, the alignment precision and signal quality indicators are given by collecting and process these information.

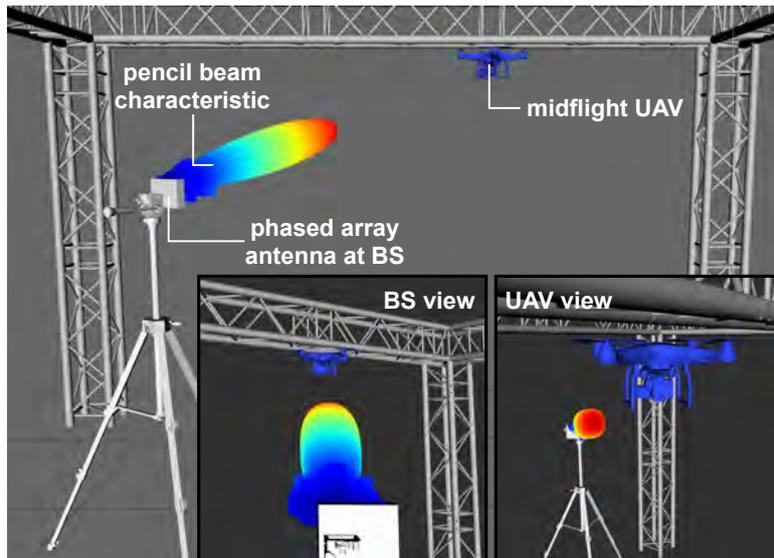


Figure 4: Illustrative virtual twin of designed testing environment

## 4 Conclusion and Further Research

The developed 5G mmWave antenna tracking evaluation system allows for connectivity measurements through indoor flight experiments as well as reproducible and controlled UAV trajectory emulation by a self-constructed rail system mounted on a truss. Experiments carried out prove, that a high performance communication link can be established and maintained by means of precise antenna beam alignment.

In future work, the testbed is planned to be used to evaluate different antenna tracking algorithms. Further, the equipment is to be extended to perform handover scenarios with the help of a second base station unit. While doing so, the beam alignment needs to be organized to seamlessly provide network coverage between two adjacent mmWave base stations.

## References

- [1] Niklas Goddemeier, Kai Daniel, and Christian Wietfeld. Role-based connectivity management with realistic air-to-ground channels for cooperative UAVs. *IEEE Journal on Selected Areas in Communications (JSAC)*, 30(5):951–963, June 2012.
- [2] Karsten Heimann, Janis Tiemann, Stefan Boecker, and Christian Wietfeld. On the potential of 5G mmwave pencil beam antennas for UAV communications: An experimental evaluation. In *WSA 2018; 22nd International ITG Workshop on Smart Antennas*, pages 1–6, March 2018.

# A Resource-Efficient Soil Moisture Sensing System

Florian Liedmann

Lehrstuhl für Kommunikationsnetze

Technische Universität Dortmund

florian.liedmann@tu-dortmund.de

Resource-efficiency is important for many different applications. Especially for a sensor network with limited energy, resources have to be used as efficient as possible. The developed Soil Moisture Sensing system uses the radio field as a sensor and provides a soil moisture estimation due to the radio field itself. By using Long Range packet radio modules, originate from the Internet of Things, a robust communication with a high receiver sensitivity can be enabled. This leads to a resource-efficient application as Wireless Underground Sensors. To ensure a minimal power consumption paired with a maximum of used communication links, a so-called rotating gateway concept was developed.

## 1 Design of the LoRa Based Soil Moisture Sensing System

In the following, the basic operating principle of the LoRa based SoMoS system will be described in detail. This is the first approach to enable a moisture sensing in an outdoor environment. The main focus is to ensure a robust sensor field with a maximal battery lifetime of the nodes, even with a low sensing resolution. Without any gateway functionality, the nodes will be called passive nodes.

LoRa packet radio is chosen as a robust communication technology with a high receiver sensitivity. LoRa provides high communication distances up to several kilometers in free space and data rates between 0.3 kbit/s and 11 kbit/s [1]. Due to the spread spectrum signal modulation technique, a signal reception even below the noise floor is possible. To maximize the communication range of the radio module, the lowest possible frequency band of 433 MHz is used. The radio module is mounted on an Adafruit Feather 32u4 development board. During active radio listening, the power consumption is about 40 mA, during transmission with +20 dBm about 120 mA and during full sleep about 0.3 mA. Due to legal requirements, the transmit power will be limited to +10 dBm (10 mW). The resulting sensor is small in size and can easily be installed with minimum

invasion. Equipped with a 3000 mAh battery, the sensor is independent of power supply. To be resistant to soil, moisture and humidity, the sensor is placed in a waterproof acrylic glass cylinder ( $\varnothing$  0.1 m) [2]. In the developed sensor field, one of the SoMoS nodes is configured as gateway to continuously listen to the communication channel and measure the RSSI of node packets. The measurement reports will then be forwarded to an embedded system with Wide Area Network (WAN) access to the monitoring system for evaluation and visualization. This transmission will be done via an Universal Serial Bus (USB) connection, which additionally avoids a runtime of only a few days of the SoMoS gateway by continuously recharging its battery. In Fig. 1 (left), the state machine of the SoMoS gateway is shown.

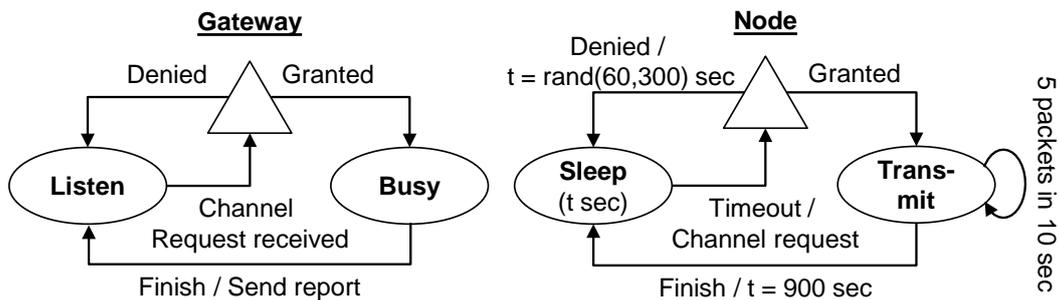


Figure 1: State machines of SoMoS gateway and passive node.

After receiving a channel request of a SoMoS node in Listen state, the node gets an acknowledgment if the channel is not occupied. Then, the gateway changes to Busy state and waits for a predefined amount of packets of the node. After a timeout, the gateway sends a report with measured and received node parameters to a monitoring system and returns to Listen state. In Fig. 1 (right), the state machine of the SoMoS node is shown. To save battery, it enters a 15 minutes Full Sleep state after successful transmission. After waking up, it requests channel access and transmits five packets in a predefined period of time if access granted. If the channel is occupied by another node, it calculates a random back-off time between 60 and 300 seconds, enters the Full Sleep state for this time and requests channel access again afterwards. The expected lifetime for both configurations is shown in Fig. 2.

Without recharging, the gateway lifetime would be 10 days only. Therefore, the gateway is connected to an embedded system via Universal Serial Bus (USB) to continuously recharge the battery and pass the report messages. The expected passive node lifetime is 397 days, which is quite enough for long-term evaluations and near to the maximum

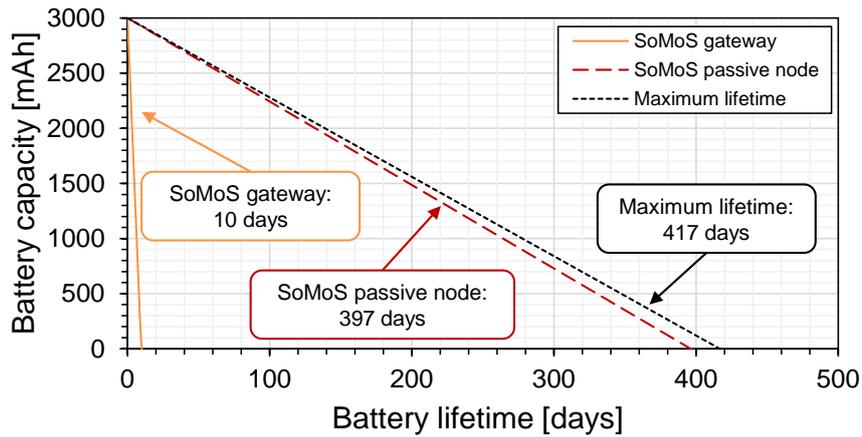


Figure 2: Expected battery lifetime of SoMoS gateway and passive node.

lifetime of 417 days, if the node would always remain in Sleep state without any packet transmissions.

## 2 Resource-Efficient Concept for Soil Moisture Sensing in the Field with Increased Resolution

The previous configuration of the SoMoS system with an USB powered gateway and passive nodes in full sleep mode enables a maximal system runtime of more than one year (see Fig. 2). But due to the star topology from nodes to gateway, the amount of communication links for the soil moisture sensing is limited. Therefore, a concept for the SoMoS system has been developed to maximize the amount of available communication links, which leads to an increased resolution of soil moisture localization. The developed, so-called rotating gateway approach increases the amount of 8 unidirectional communication links to 36 bidirectional communication links with 72 measurable RSSI values (see Fig. 3).

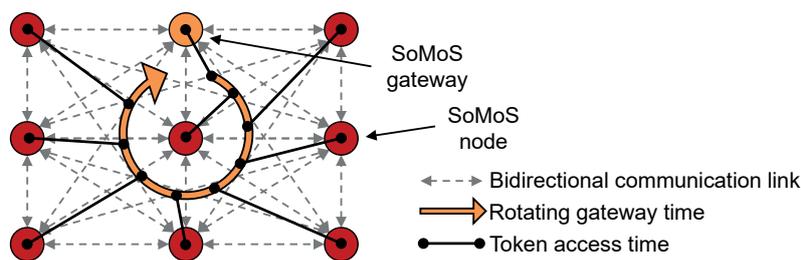


Figure 3: Rotating gateway concept to increase soil moisture detection resolution in the field [3].

The main challenge of this configuration is to minimize the clock drift caused by the local oscillators of each node, which is constant within one system, but differs from the other clock drifts. Especially during the full sleep mode, most components for a better clock synchronization are switched off. For a target sleep time of 900 s, this drift results

in a maximal observed time deviation of up to 70 s, in which the node wakes up too early or too late. Due to the fact that all nodes have to be active, this would result in an significantly decreased battery lifetime. Therefore, each node calibrates itself to the time of the first token sent by the gateway. If a node misses the first token and wakes up later, it will ignore the sleep request and remain awake waiting for the next token [3].

In this configuration with active nodes, the battery lifetime obviously is shorter than before with passive nodes, but 360 days can still be expected (see Fig. ??).

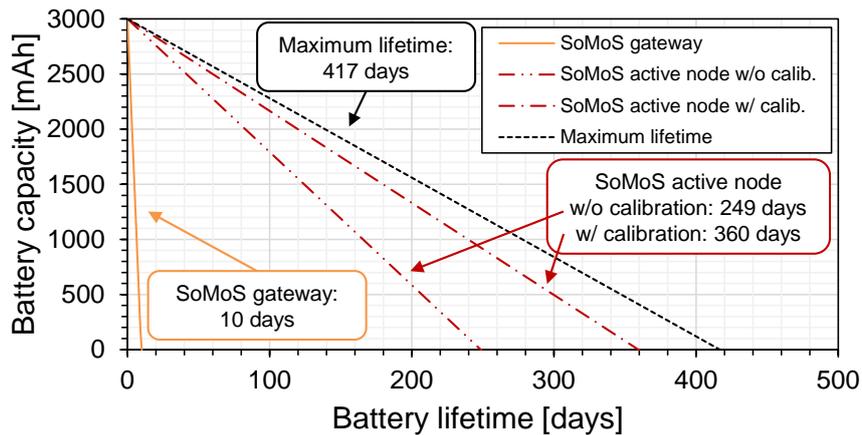


Figure 4: SoMoS battery consumption with maximized link amount.

### 3 Conclusion and Outlook

This contribution shows the feasibility of a multidimensional radio field based Soil Moisture Sensing System with a resource-efficient concept for moisture sensing in the field with increased resolution.

In future work, the results of the SoMoS system can be evaluated in machine-learning based approaches to increase the efficiency of soil moisture detection.

### References

- [1] P. Jörke, S. Böcker, F. Liedmann, and C. Wietfeld, "Urban Channel Models for Smart City IoT-Networks Based on Empirical Measurements of LoRa-links at 433 and 868 MHz," in *IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC) - Workshop WS-01 on "Communications for Networked Smart Cities (CorNer)*, Oct 2017.
- [2] F. Liedmann and C. Wietfeld, "Somos - a multidimensional radio field based soil moisture sensing system," in *IEEE Sensors*, (Glasgow, Scotland, UK), oct 2017.
- [3] F. Liedmann, C. Holewa, and C. Wietfeld, "The radio field as a sensor—a segmentation based soil moisture sensing approach," in *Sensors Applications Symposium (SAS), 2018 IEEE*, pp. 1–6, IEEE, 2018.

# Opportunistic Transmission of Vehicular Sensor Data Leveraging Machine Learning-based Data Rate Prediction

Benjamin Sliwa  
Lehrstuhl für Kommunikationsnetze  
Technische Universität Dortmund  
benjamin.sliwa@tu-dortmund.de

Exploiting cars as moving sensor nodes is an enabler for data-driven crowd-sensing services in Intelligent Transportation Systems (ITSs). However, the massive increases in cellular data transmissions lead to resource competition among the different cell users. Since straightforward (periodic) transmission approaches do not consider the network quality within the data transfer process, the resource-efficiency is low as transmissions are performed during low network quality periods, which require retransmissions and long transmission durations. In this report, we summarize our work on opportunistic data transfer that exploits machine learning-based data rate prediction to schedule sensor data transmissions in a context-aware manner. The real-world experiments show that the proposed approach is able to achieve massive increases in the end-to-end data rate and additionally lowers the average power consumption of the mobile device.

## 1 Solution Approach

The general approach of the proposed opportunistic transmission scheme is to exploit *connectivity hotspots*, where reliable and resource-efficient data transfer is expected and

to avoid transmissions during *connectivity valleys*. For achieving this behavior, a probabilistic process is applied that determines a transmission probability based on the predicted data rate. Transmissions are delayed if the expected data rate is low and performed early if the anticipated throughput is high.

Based on the Signal-to-interference-plus-noise Ratio (SINR)-based Channel-aware Transmission (CAT) scheme, we propose a machine learning-based mechanism called Machine Learning CAT (ML-CAT) [2]. For predicting the achievable data rate  $\tilde{S}(t)$ , the M5 Regression Tree (M5T) model is applied based on the downlink indicators Reference Signal Received Power (RSRP), Reference Signal Received Quality (RSRQ), SINR, Channel Quality Indicator (CQI), the vehicle's velocity and the payload size of the data packet. The transmission probability  $\rho(t)$  is then computed as

$$\rho(t) = \left( \frac{\tilde{S}(t)}{S_{\max}} \right)^{\alpha} \quad (1)$$

with respect to a reference maximum  $S_{\max}$  and  $\alpha$  being an exponent to control the preference of high metric values.

With Machine Learning Predictive CAT (ML-pCAT) [3], the transmission scheme is extended by a predictive component. Mobility prediction that exploits navigation system knowledge is applied to forecast the future location  $\tilde{\mathbf{P}}(t + \tau)$  for a defined prediction horizon  $\tau$ . Crowdsensing-based connectivity maps are used as a priori information that allow to derive channel quality estimations based on location information. The predicted position is used to look up the network quality context  $\tilde{\mathbf{C}}(t + \tau)$ , which is then utilized to predict the anticipated future data  $\tilde{S}(t + \tau)$  rate using M5T. ML-pCAT then computes the transmission probability with respect to the current and the expected future data rate as

$$\rho(t) = \left( \frac{\tilde{S}(t)}{S_{\max}} \right)^{\alpha \cdot z} \quad (2)$$

with

$$z = \begin{cases} \max(|\Delta S(t) \cdot (1 - S(t)) \cdot \beta|, 1) & : \Delta S(t) > 0 \\ (\max(|\Delta S(t) \cdot S(t) \cdot \beta|, 1))^{-1} & : \Delta S(t) \leq 0 \end{cases} \quad (3)$$

and  $\beta$  being an additional parameter for controlling the impact of the prediction for the overall transmission probability.

## 2 Real World Performance Evaluation

The performance of the proposed transmission scheme is evaluated in the public cellular network on different tracks (suburban and highway traffic characteristics). For the data transfer, an Android-based User Equipment (UE) is used, which executes the ML-CAT

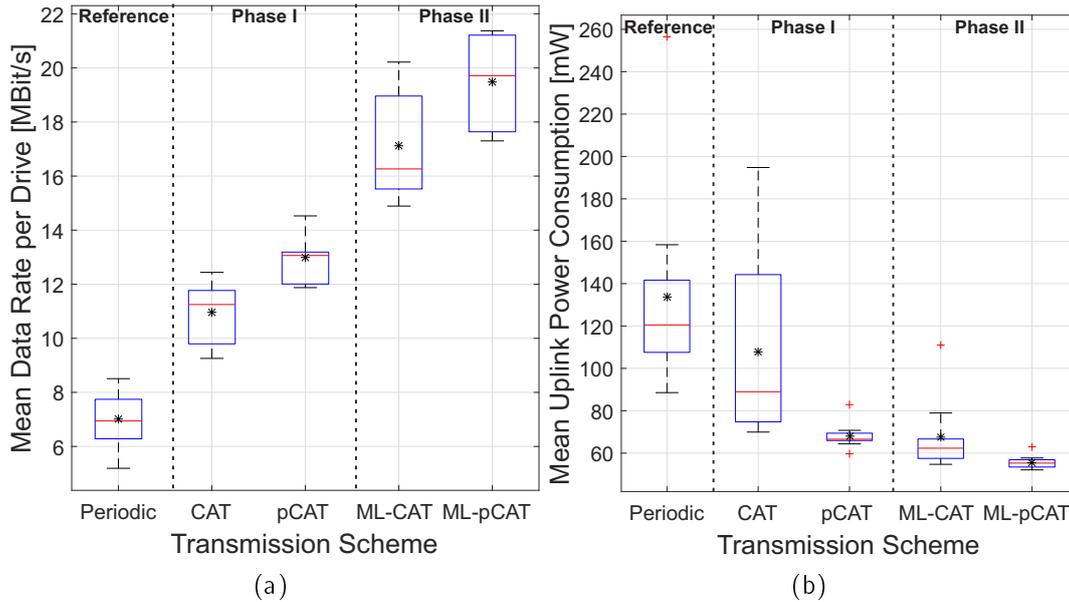


Figure 1: Results of the real-world evaluation campaign. ML-CAT and ML-pCAT significantly outperform the straightforward approach as well as the SINR-based CAT variants.

application and transmits the data of a virtual sensor application to a cloud-based measurement server using the Long Term Evolution (LTE) network. For evaluating the energy-efficiency of the proposed approach, we apply the Context-aware Power Consumption Model (CoPoMo) [1], which is extended by a machine learning-based mechanism to derive the applied transmission power based on the measured passive downlink indicators.

Fig. 1 shows the results for data rate and power consumption. The results of straightforward periodic data transmission and the SINR-based CAT groundwork are provided for comparison. While the periodic approach does not consider the network quality at all, the SINR-based variants of CAT are able to increase the average throughput by applying a channel-aware transmission scheduling scheme. However, it can be seen that the proposed approach leads to massive increases in the achieved data rate. The predicted throughput provides a better metric for the channel quality than just SINR measurements as the integration of the payload size immanently considers the dependency of transmission duration and channel coherence time. In addition, protocol-related effects such as the slow start of Transmission Control Protocol (TCP) are also contained in the prediction.

As a side-effect, also the power consumption of the mobile UE is reduced significantly. While one reason for this behavior is the reduced average transmission power due to the exploitation of *connectivity valleys*, in addition the actual transmissions are performed

faster due to the higher data rate. Therefore, the UE is able to return to the low power IDLE state earlier.

### 3 Conclusion and Further Research

With ML-CAT, we have demonstrated the potentials of using machine learning-based data rate prediction as a metric to perform vehicular sensor data transfer in a context-aware manner. In contrast to existing network quality indicators, hidden interdependencies between the application, the protocols and the channel quality are implicitly considered. In future work, we will investigate multi-connectivity approaches that exploit the joint-usage of different communication networks.

### References

- [1] Robert Falkenberg, Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. Machine learning based uplink transmission power prediction for LTE and upcoming 5G networks using passive downlink indicators. In *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, Chicago, USA, Aug 2018.
- [2] Benjamin Sliwa, Thomas Liebig, Robert Falkenberg, Johannes Pillmann, and Christian Wietfeld. Efficient machine-type communication using multi-metric context-awareness for cars used as mobile sensors in upcoming 5G networks. In *2018 IEEE 87th Vehicular Technology Conference (VTC-Spring)*, Porto, Portugal, Jun 2018. Best Student Paper Award.
- [3] Benjamin Sliwa, Thomas Liebig, Robert Falkenberg, Johannes Pillmann, and Christian Wietfeld. Machine learning based context-predictive car-to-cloud communication using multi-layer connectivity maps for upcoming 5G networks. In *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, Chicago, USA, Aug 2018.

# Application of dynamic network options to minimize traffic congestions in inner-city networks

Tim Vranken  
Physik von Transport und Verkehr  
Universität Duisburg-Essen  
tim.vranken@uni-due.de

This report sums up our recent research on applying the newly introduced trip planning method [7] (following Combination)— as well as the "breakdown minimization principle" [2] (following BMP) and "Wardrop's user Equilibrium" [3] (following WE)— together with dynamic network options to avoid traffic congestions in inner city systems. To this end, we simulated multiple different situations, which can occur in a network, and tested their impacts on the global travel time.

## 1 Introduction

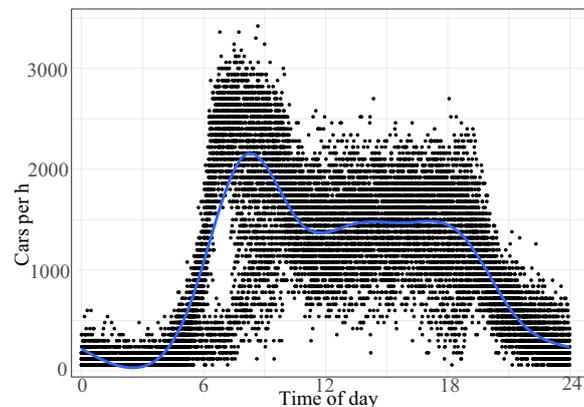
Traffic volume in inner cities is often at, or even above, the road capacity which results in a lot of congestion and increased travel times. Since the space in cities is limited, this problem can not be fixed by a static increase of capacity (e.g. additional lanes or streets). This leaves two solutions, either the traffic capacity has to be optimized in such a way that it's better fitting to the traffic volumes, or the traffic volume has to decrease. While the former can be attempted by dynamic traffic options as proposed in [7], the latter can hopefully be solved by automated vehicle.

## 2 Different Dynamic Situations

To test the Impact of dynamic network options, a simulation of the Düsseldorf inner city has been created (see fig. 1(a)). In order to recreate the network in a realistic way, traffic flow data from a video detection system, as well as from floating car data is taken (see fig. 1(b) ) and the traffic volume is designed in accordance with the blue line.



(a) Marked area in the Duesseldorf inner city. The blue marked streets and intersections are streets used in the simulation. The green marked streets act as sources and the red marked streets are sinks. The streets marked with a red x have a road closed in a later simulation. The intersections marked with a green circle have dynamic lanes that change in an additional simulation based on the time of the day to adjust to traffic demand. (Map: ©OpenStreetMap contributors, CC BY-SA)



(b) Average cars per hour at different times of the day for a street with two lanes in Duesseldorf over two month obtained from real-world detectors. The Blue line marks the average traffic flow at the time.

In this network three different simulations were done:

1. 100 normal days without changes done to the system were simulated.
1. 100 days, where a street was closed at 6.30 o'clock were simulated.
2. 100 days with a dynamic lane were simulated.

Table 1: Average travel time  $\bar{\tau}$  and the standard error  $se$  for all four methods and the three simulations.

Method	$\bar{\tau}$ normal	$se$ normal	$\bar{\tau}$ Dynamic	$se$ Dynamic	$\bar{\tau}$ Closing	$se$ Closing
WE	589.2	1.9	545.1	1.8	568.96	2
BMP	449.7	0.3	448.2	0.3	448.9	0.3
Combination	408.9	0.4	409.2	0.4	410	0.5

The dynamic lane leads into the city from 0-12 o'clock and out of the city from then on till midnight. Because Düsseldorf is a big commuter city, more people drive into the city to work than drive out of it. The dynamic lane should thus help by increasing the traffic capacity leading into the city in the morning and out of the city in the afternoon without the need to build a additional lane. The results of the simulations can be seen in table 1. One can see, that neither closing a street nor using a lane as a dynamical lane has a significant impact on the BMP and the Combination. The dynamic lane results in no improvement because the additional traffic flow that passes the bottleneck can not be supported by the following bottlenecks, so that the breakdown is only shifted in stead of prevented. Surprisingly the simulations with a closed lane resulted in lower average travel times than the normal simulation for WE. In the normal simulation the closed road is the shores way for a lot of commuters to get into the city thus a lot of them choose this road which results in large congestion and longer travel times. Since there are more than one alternative roads which are only slightly longer, the vehicle distribute them on these roads, which reduces the average congestion on the roads and thus the travel times. This shows clearly that the average use of the complete traffic capacity of cities is not well used by commuters.

### 3 Realistic inner city traffic

The results presented in the previous chapter were obtained by applying a modified version of the break-light-Model [4]. In order to create even more realistic traffic, currently a inner city Simulation-Model based on the Pottmeier-Model [6] —which is a accident free version of the Lee-Model [5]— is being created and tested. This model has two important improvements to the former model. First it has a fixed maximal deceleration of  $D = 2$  which means vehicle have to react pro active on the surrounding traffic to prevent accidents. Second it introduces driver behaviour. Drivers can behave passive or aggressive in traffic, which means they can, if they judge the situation save, drive closer to the vehicle in front than the minimal safety distance. Thorough this Agents have to react stronger if the car in front of them dawdle which results in more realistic jam pattern. This will also help to introduce and differentiate automated vehicle from human driven once.

## 4 Automated vehicle

As described in the previous chapter, in the following work, automated vehicle will be introduced in the model and their impact on the traffic flow will be tested. Since automated vehicle always have to strictly follow the rules, it is to be expected, that they will first reduce traffic flow in mixed traffic. For example, they can't reduce their distance to the vehicle in front below the safety distance which means than in synchronized flow, the average distance between vehicles increases and thus the maximum flow decreases at the same average speed than in pure human driven traffic. In resent theory's [8] it is even proposed that the gap between automated and human driven vehicles has to be further increased, which would mean even lower traffic volumes. Thus a analysis of mixed traffic flow and how to fix the coming problems will be done.

## References

- [1] B. S. Kerner. Optimum principle for a vehicular traffic network: minimum probability of congestion. *Journal of Physics A: Mathematical and Theoretical*, 44(9):092001, 2011.
- [2] B. S. Kerner. Breakdown minimization principle versus wardrop's equilibria for dynamic traffic assignment and control in traffic and transportation networks: A critical mini-review. *Physica A: Statistical Mechanics and its Applications*, 466, 09 2016.
- [3] W. Knospe, L. Santen, A. Schadschneider, and M. Schreckenberg. Towards a realistic microscopic description of highway traffic. *Journal of Physics A: Mathematical and General*, 33(48):477–485, 2000.
- [4] Hyun Keun Lee, Robert Barlovic, Michael Schreckenberg, and Doochul Kim. Mechanical restriction versus human overreaction triggering congested traffic states. *Phys. Rev. Lett.*, 92:238702, Jun 2004.
- [5] Andreas Pottmeier. *Realistic Cellular Automaton Model for Synchronized Two-Lane Traffic*. PhD thesis, Universtität Duisburg-Essen, 2007.
- [6] Tim Vranken, Benjamin Sliwa, Christian Wietfeld, and Michael Schreckenberg. Performance comparison of dynamic vehicle routing methods for minimizing the global dwell time in upcoming smart cities. In *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, Chicago, USA, Aug 2018.
- [7] Lanhang Ye and Toshiyuki Yamamoto. Modeling connected and autonomous vehicles in heterogeneous traffic flow. *Physica A: Statistical Mechanics and its Applications*, 490, 08 2017.





## Subproject C1

Feature selection in high dimensional data  
for risk prognosis in oncology

Sven Rahmann

Alexander Schramm

# Controlling the False Discovery Rate in Boolean Matrix Factorization

Sibylle Hess

Lehrstuhl für Künstliche Intelligenz

Technische Universität Dortmund

sibylle.hess@tu-dortmund.de

When it comes to clustering nonconvex shapes, two paradigms are used to find the most suitable clustering: minimum cut and maximum density. The most popular algorithms incorporating these paradigms are Spectral Clustering and DBSCAN. Both paradigms have their pros and cons. While minimum cut clusterings are sensitive to noise, density-based clusterings have trouble handling clusters with varying densities. In this paper, we propose SpectACI: a method combining the advantages of both approaches, while solving the two mentioned drawbacks. Our method is easy to implement, such as spectral clustering, and theoretically founded to optimize a proposed density criterion of clusterings. Through experiments on synthetic and real-world data, we demonstrate that our approach provides robust and reliable clusterings.

**Introduction** Despite being one of the core tasks of data mining, and despite having been around since the 1930s, the question of clustering has not yet been answered in a manner that doesn't come with innate disadvantages. The report that you are currently reading will also not provide such an answer. Several advanced solutions to the clustering problem have become quite famous, and justly so, for delivering insight in data where the clusters do not offer themselves up easily. *Spectral Clustering* provides an answer to the curse of dimensionality inherent in the clustering task formulation, by reducing dimensionality through the spectrum of the similarity matrix of the data. *DBSCAN* is a density-based clustering algorithm which has won the SIGKDD test of time award in 2014. Both Spectral Clustering and DBSCAN can find non-linearly separable clusters, which trips up naive clustering approaches; these algorithms deliver good results. In this paper, we propose a new clustering model which encompasses the strengths of both

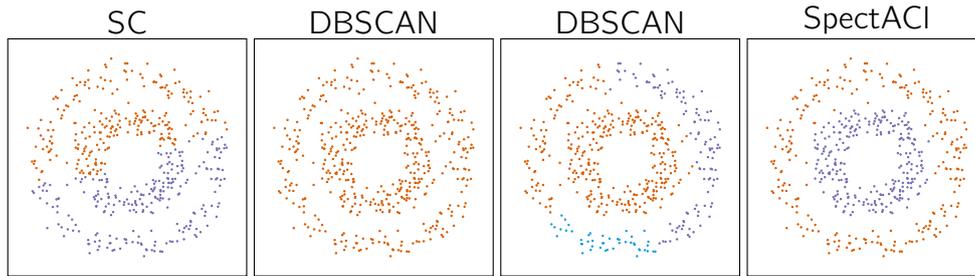


Figure 1: Performance of Spectral Clustering, DBSCAN (using  $minPts = 25$  and  $26$ ), and SpectACI on two concentric circles. Best viewed in color.

Spectral Clustering and DBSCAN; the combination can overcome some of the innate disadvantages of both individual methods.

For all their strengths, even the most advanced clustering methods nowadays still can be tripped up by some pathological cases: datasets where the human observer immediately sees what is going on, but which prove to remain tricky for all state-of-the-art clustering algorithms. One such example is the dataset illustrated in Figure 1: we will refer to it as the *two circles* dataset. As the leftmost plot in Figure 1 shows, Spectral Clustering does not at all uncover the innate structure of the data. It is well-known that Spectral Clustering is highly sensitive to noise; if two densely connected communities are additionally connected to each other via a narrow bridge of only a few observations, spectral clustering runs the risk of reporting these communities plus the bridge as a single cluster, whereas two clusters plus a few noise observations (outliers) would be the desired outcome. The middle plots in Figure 1 shows how DBSCAN (with  $minPts$  set to 25 and 26, respectively) fails to realize that there are two clusters. This is hardly surprising, since DBSCAN is known to struggle with several clusters of varying density, and that is exactly what we are dealing with here: since both circles consist of the same number of observations, the inner circle is substantially more dense than the outer circle. The rightmost plot in Figure 1 displays the result of the new clustering method that we introduce in [1], SpectACI (*Spectral Averagely-dense Clustering*): it accurately delivers the clustering that represents the underlying phenomena in the dataset.

SpectACI combines the benefits of Spectral Clustering and DBSCAN, while alleviating some of the innate disadvantages of each individual method. Our method finds clusters having a large average density, where the appropriate density for each cluster is automatically determined through the spectrum of the weighted adjacency matrix. Hence, SpectACI does not suffer from sensitivity to the  $minPts$  parameter as DBSCAN does, and unlike DBSCAN it can natively handle several clusters with varying densities. As in the Spectral Clustering pipeline, the final step of SpectACI is an embedding postprocessing step using  $k$ -means. However, unlike Spectral Clustering, we demonstrate the fundamental soundness of applying  $k$ -means to the embedding step in SpectACI: from

SpectACI's objective function we derive an upper bound by means of the eigenvector decomposition; we derive that the optimization of our upper bound is equal to  $k$ -means on the eigenvectors. Our Python implementation, and the data generating and evaluation script, are publicly available<sup>1</sup>.

**Spectral Averagely-Dense Clustering** We propose a cluster definition based on the average density (i.e., node degree) in the subgraph induced by the cluster. Let  $W$  be the adjacency matrix. Since we are interested in deviations of the nodes' degree, we employ the  $\epsilon$ -neighborhood graph to determine  $W$ . We strive to solve the following problem, maximizing the *average cluster density*:

$$\max_{Y \in \mathbb{1}^{m \times r}} \text{tr}(Y^T W Y (Y^T Y)^{-1}) = \sum_s \delta(Y_{\cdot s}, W), \quad (1)$$

where  $\mathbb{1}^{m \times r}$  denotes the space of binary  $m \times r$  matrices which represent a partition of the data points, that is  $|Y_{j \cdot}| = 1$ . The objective function returns the sum of average node degrees  $\delta(Y_{\cdot s}, W)$  in the subgraph induced by cluster  $s$ , i.e.:

$$\delta(Y_{\cdot s}, W) = \frac{Y_{\cdot s}^T W Y_{\cdot s}}{\|Y_{\cdot s}\|^2} = \frac{1}{|Y_{\cdot s}|} \sum_{j: Y_{js}=1} W_{j \cdot} Y_{\cdot s}. \quad (2)$$

The Objective (1) is equivalent to minimum cut if the matrix  $W$  is normalized. We derive a new relationship for the solution of Objective (1) and the spectrum of the adjacency matrix. Thereby, we establish a connection with the application of  $k$ -means to the spectral embedding. As a result, our method encompasses the same steps as Spectral Clustering, and hence it can be efficiently computed even for large scale data.

The function  $\delta$  from Eq. (2) is also known as the Rayleigh quotient. The values of the Rayleigh quotient depend on spectral properties of the applied matrix. A simple calculation shows that the eigenvectors  $V_1, \dots, V_d$  to the  $d$ -largest eigenvalues of  $W$  span a space whose points have a minimum density  $\delta(y, W) \geq \lambda_d$ . Thus, the projection of  $W$  onto the subspace spanned by the first eigenvectors reduces the dimensionality of the space in which we have to search for optimal clusters.

This insight suggests a naive approach, where we approximate  $W$  in Objective (1) by its truncated eigendecomposition. This results in a trace maximization problem which is equivalent to  $k$ -means clustering on the data matrix  $U = V^{(r)} (\Lambda^{(r)})^{1/2}$ . Unfortunately, this approach does not yield acceptable clusterings. The objective of  $k$ -means is nonconvex, having multiple local solutions. The number of local solutions increases with every eigenvector which we include in the eigendecomposition from  $W$ , due to the orthogonality of eigenvectors. As a result,  $k$ -means returns clusterings whose objective function value

---

<sup>1</sup><https://sfb876.tu-dortmund.de/spectacl>

approximates the global minimum, but which reflect seemingly haphazard groupings of the data points.

The eigenvectors of  $W$  are not only orthogonal, but also real-valued, having positive as well as negative values. The first eigenvectors have a high value of the density function  $\delta$ , but the mixture of signs in its entries makes an interpretation as cluster indicators difficult. Aiming for an interpretable embedding, we make the following observation.

**Observation 1.** *Let  $W$  be a symmetric real-valued matrix, and let  $v$  be an eigenvector to the eigenvalue  $\lambda$ . Let  $v = v^+ - v^-$ , with  $v^+, v^- \in \mathbb{R}_+^m$  be the decomposition of  $v$  into its positive and negative parts. The nonnegative vector  $u = v^+ + v^-$  has a density*

$$\delta(u) \geq |\lambda|.$$

We refer to the eigenvectors, whose entries are replaced with their absolute values, i.e.,  $u_j = |v_j|$  for  $1 \leq j \leq m$ , as *projected eigenvectors*. The projected eigenvectors have locally similar values and an interpretation as fuzzy cluster indicators. Furthermore, projected eigenvectors are not orthogonal and provide multiple views on possible dense and fuzzy clusters, which is beneficial in order to robustly identify the clustering structure. We summarize the resulting method SpectACI (Spectral Averagely-dense Clustering) with the following steps:

1. compute the adjacency matrix  $W$ ;
2. compute the truncated eigendecomposition  $W \approx V^{(d)} \Lambda^{(d)} V^{(d)\top}$ ;
3. compute the projected embedding  $U_{jk} = |V_{jk}^{(d)}| |\lambda_k|^{1/2}$ ;
4. compute a  $k$ -means clustering, finding  $r$  clusters on the embedded data  $U$ .

The adjacency matrix  $W$  is here calculated by the  $\epsilon$ -neighborhood graph. However, our method is in principle applicable to any provided (weighted) adjacency matrix. We recall that the average density objective is equivalent to the ratio cut objective if we normalize the matrix  $W$ . We refer to this version as normalized SpectACI. In the normalized case, we compute the adjacency matrix according to the  $k$ -nearest neighbor graph, which is suggested for applications of Spectral Clustering. Through experiments on synthetic and real-world data, we demonstrate that our approach provides robust and reliable clusterings.

## References

- [1] Sibylle Hess, Wouter Duivesteijn, Philipp Honysz, and Katharina Morik. The SpectACI of nonconvex clustering: A spectral approach to density-based clustering. In *AAAI*, 2019.

# **Evaluation of Protein Kinase C Iota (PRKCI) as a target for the treatment of recurrent neuroblastoma**

Marc Schulte

Molecular Oncology, University Hospital Essen

marc.schulte@uk-essen.de

Neuroblastoma (NB) is a common childhood cancer most often situated in the adrenal glands, but can also develop in the neck, chest, abdomen, or spine. Even if the tumor shows good response to an initial treatment, difficult to treat metastases and recurrent tumors are often developed. We previously identified several potential key genes involved in NB relapse by analyzing genome sequencing and expression data of paired primary and recurrent NB. One of these genes is Protein Kinase C Iota (PRKCI) for which we observed an increased expression of in recurrent NB compared to the paired primary tumors. In order to validate these data in a biological system, we use CRISPR / Cas9 technology to study the function of PRKCI in the NB cell line SH-EP. We established CRISPR / Cas9 –mediated overexpression and knockout of PRKCI and analyzed its role on proliferation, clonal outgrowth, migratory capacity, invasiveness as well as on drug resistance. Here we show that expression PRKCI is increased by overexpression of MYCN, a known oncogene in NB. Further, we were able to show in 3D culture experiments that PRKCI modulates invasion of tumor cells in the surrounding matrix and thus may have an important function in relapse and metastasis formation.

Neuroblastoma is the most common extra cranial solid tumor in childhood and accounts for 7 -10 % of all childhood cancers. As a tumor of the autonomic nervous system, NB derives from neural crest tissue and thus usually arises in a paraspinal location in the abdomen or chest [2]. Thanks to improved therapies, NB often initially responds very well to the treatment. However, at relapse there is only very little to offer for the patients and hence relapses correlate with poor prognosis and fatal outcome. Previously, we used whole-exome sequencing, mRNA expression profiling, array CGH and DNA methylation

analyses to characterize 16 paired samples at diagnosis and relapse from individuals with NB [4]. We observed that Protein Kinase C Iota (PRKCI) was significantly higher expressed in recurrent NB compared to the corresponding primary tumors (Figure 1a) [5]. Moreover, an elevated PRKCI expression correlates with NB relapse and poor survival (Figure 1b).

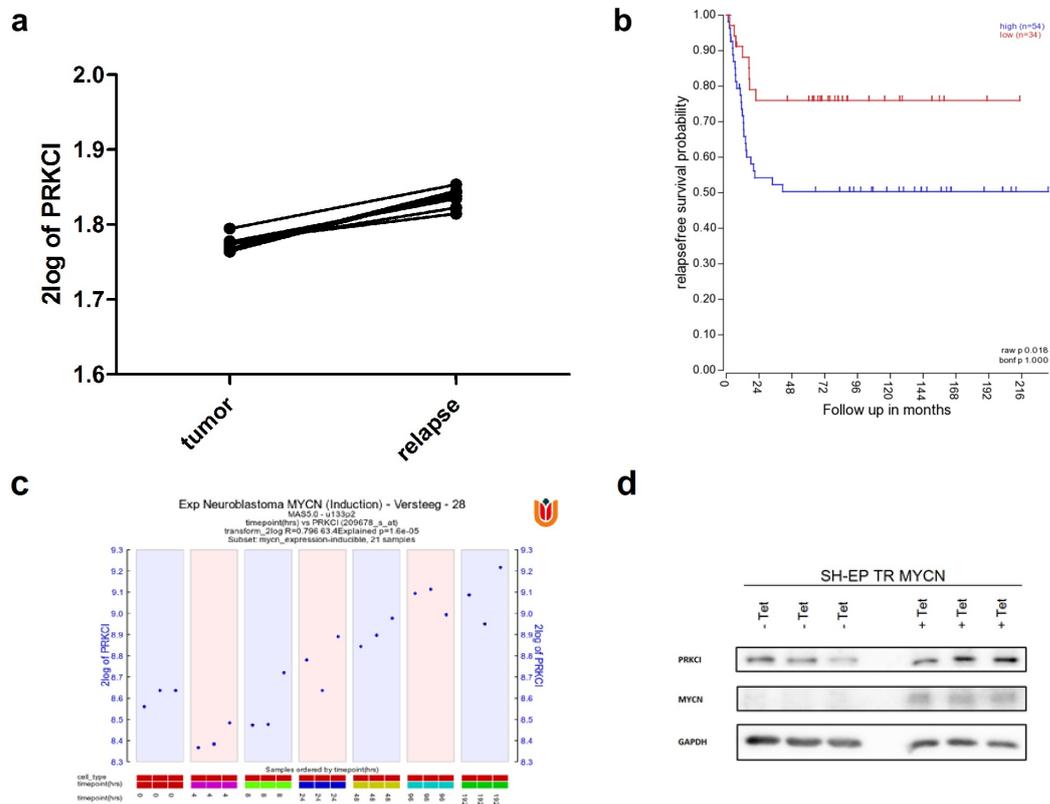


Figure 1: **PRKCI expression is increased in relapsing NBs and high PRKCI levels can be achieved by overexpression of MYCN and correlated with unfavorable clinical outcome (R2: Genomics Analysis and Visualization Platform, AMC Amsterdam).** (a) PRKCI expression level in primary and relapse NBs (data from [5]). (b) Kaplan curve of relapse-free survival for patients with PRKCI low and high NBs (data from Valentijn et al., PNAS [6]). (c) Overexpression of MYCN in the MYCN inducible NB cell line SK-N-AS results in increased PRKCI RNA levels (data from Koppen et al., Int J Cancer [1]). (d) Overexpression of MYCN in the MYCN inducible NB cell line SH-EP results in increased PRKCI protein 72h after start of induction.

Data from Bandino et al. show that overexpression of the known oncogene MYCN results in increased expression of PRKCI suggesting that the PRKCI gene might be a downstream target of MYCN (Figure 1c)[1]. By use of the MYCN inducible NB cell line

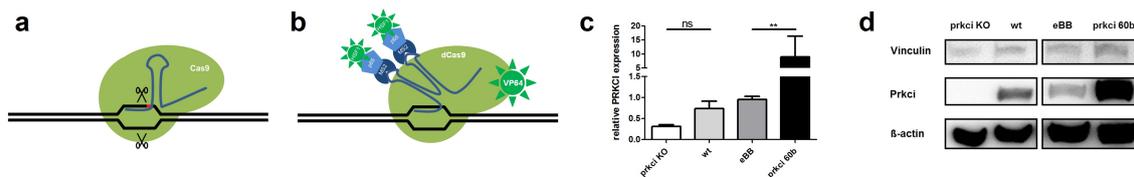


Figure 2: **CRISPR / Cas9 based-approaches to knockout and overexpress PRKCI in NB cells.** (a) Schematic representation of double strand break generation by Cas9 to create a gene knock out (KO). (b) Schematic representation of CRISPR SAM [3], which requires incorporation of three activation domains, VP64, P65 and HSF1, into the nuclease inactive Cas9 complex to activate transcription. (c) Confirmation of PRKCI KO and overexpression in SHEP cells (KO: knockout, wt: wild type, eBB: empty backbone control, 60b: overexpression) by real time PCR. (d) Western Blot analysis confirming KO and overexpression, respectively, of PRKCI in SHEP cells.

SH-EP, we were able to confirm this observation on RNA (data not shown) and protein level (Figure 1d).

We further used different CRISPR / Cas9 based-approaches to knockout (KO) or overexpress PRKCI in SH-EP cells. While the KO was based on a double strand break induced by wild-type Cas9 (Figure 2a)[4], the CRISPR/Cas9 Synergistic Activation Mediator (SAM) system was used to overexpress PRKCI (Figure 2b) [3]. Both, knockout and overexpression were confirmed by quantitative RT-PCR (Figure 2c) and Western Blot analysis (Figure 2d).

Our results revealed that PRKCI knock out surprisingly increased proliferation compared to parental cells (Figure 3a). Moreover, the migratory capacity was reduced for PRKCI KO cells compared to the parental cells (Figure 3b). Based on this observation, we further investigated the invasive capacity of the cells in an 3D-spheroid assay. As shown in Figure 3c, knock out of PRKCI results in a decreased spheroid formation ability and invasiveness, while overexpression of PRKCI increases the invasiveness of SH-EP cells.

Although PRKCI is found to be elevated in aggressive and recurrent NB tumors, both CRISPR-mediated PRKCI knockout or overexpression resulted in formation of viable subclones in SH-EP cells. While PRKCI knockout reduced migratory capacity, it surprisingly increased proliferation. Further investigation of cell-cell adhesion and invasiveness revealed that knockout of the *prkci* gene impaired cell aggregation and reduced the invasiveness of SH-EP cells in our 3D model. Moreover, the invasiveness of PRKCI overexpressing SH-EP cells is increased compared to the empty backbone control.

Thus, PRKCI plays an important role in cellular adhesion and the processes regulating SH-EP cell spheroid growth. These findings may also apply to NB in general, but in vivo experiments will be necessary to clarify the role of PRKCI in modulating NB aggressiveness.

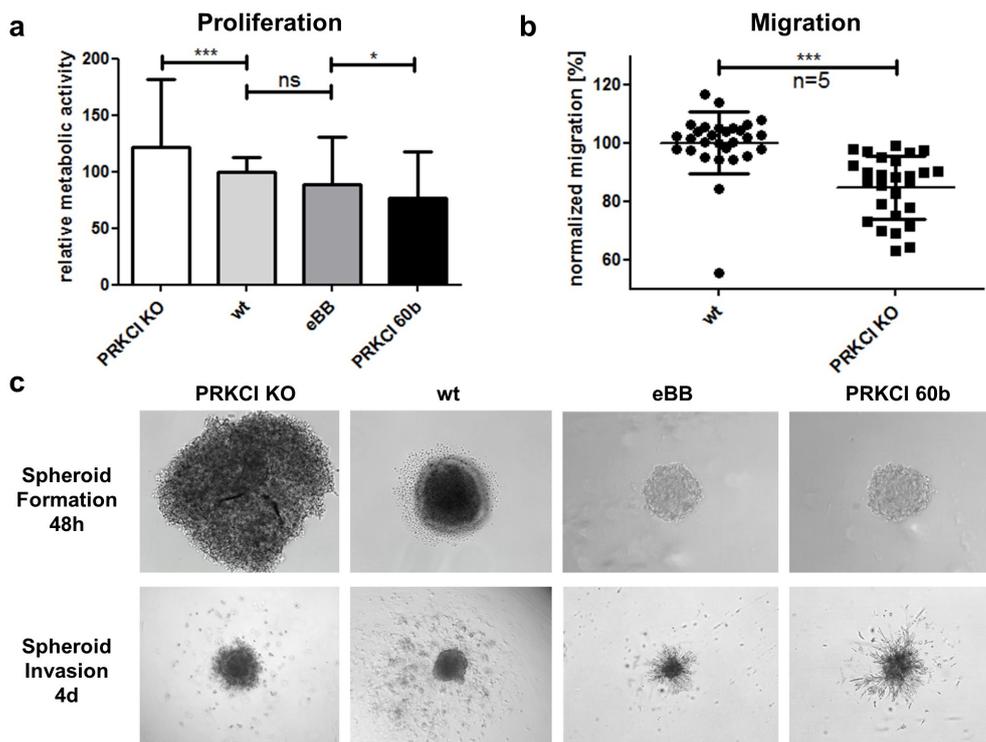


Figure 3: **Proliferation, migratory capacity and invasiveness of SH-EP cells with up- or downregulated PRKCI expression in comparison to the parental cells and empty backbone control (wt, eBB).** (a) 10000 cells/well were incubated for 48h and metabolic activity was determined. (b) Cells at confluency were scratched and migration into the scratch was documented at 0h and 24h (evaluation by Tscratch software; Gebäck and Schulz, ETH Zürich) (c) 2000 cells were seeded in plates with cell repellent surface and imaged after 48h. After complete spheroid formation, spheroids were embedded in Matrigel and images were taken four days later. Representative pictures are shown.

[1] Koppen, A et al. (2008) Dickkopf-3 expression is a marker for neuroblastic tumor maturation and is down-regulated by MYCN. *International journal of cancer*, 122(7):1455-64.

[2] Brodeur, G. (2003). Neuroblastoma: biological insights into a clinical enigma. *Nature Reviews Cancer*, 3(3), pp.203-216.

[3] Konermann, S et al. (2015) Genome-scale transcriptional activation by an engineered CRISPR-Cas9 complex. *Nature*, 517(7536):583-8.

[4] Ran, F et al. (2013). Genome engineering using the CRISPR-Cas9 system. *Nat Protoc*, 8(11), pp.2281-2308.

[5] Schramm, A et al. (2015). Mutational dynamics between primary and relapse neuroblastomas. *Nature Genetics*, 47(8), pp.872-877.

[6] Valentijn, LJ et al. (2012) Functional MYCN signature predicts outcome of neuroblastoma irrespective of MYCN amplification. *Proc Natl Acad Sci USA*, 109(47):19190-5.

# Properties of Winnowed MinHashing

Henning Timm

Genome Informatics, Institute of Human Genetics  
University Hospital Essen, University of Duisburg-Essen  
henning.timm@tu-dortmund.de

31.12.2018

Continuing from the development of an index data structure for metagenomic screening, which relies on the decomposition of the reference sequence into segments of variable length, we analyzed the distribution of segment lengths to better predict index sizes. We were able to identify a characteristic distribution of segment length and find a recursive formula to model this distribution.

## Introduction and Problem Definition

A core question in metagenomics is the classification of potentially unknown organisms in a biological sample. This can be achieved by DNA sequencing and a genomic reference database.

Starting from the biological sample, for example a scoop of water taken from a pond, a sequencing experiment is performed. This sample contains an unknown number of individuals of different, mostly unknown, species. Since for most species no reference genome is available, protein databases are used to classify the species present in the sample. These databases contain proteins that are either associated with a specific biological function, for example photosynthesis, or with a specific taxonomic unit. If a read from the sample maps to a protein in the database, this provides evidence that a species that is associated with the protein was present in the sample.

In the context of computer science, metagenomic screening is a text search problem. Typically these analyses are performed using  $q$ -grams, i.e. overlapping substrings of length  $q$  of the texts. For a small document  $R$  (a read), the most similar protein sequences

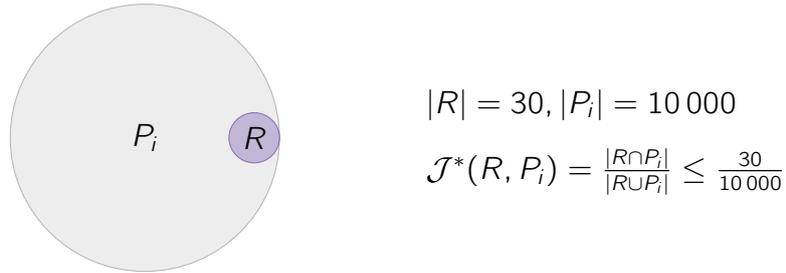


Figure 1: If the size of the read  $R$  and the protein  $P_i$  differ, the maximal possible resemblance is limited.

have to be found in a large reference database  $P = \{P_i\}_{i=1}^m$ . To avoid the pairwise comparison of  $R$  with all elements of  $P$ , the proteins can be indexed using MinHashing to identify similar pairs in linear time.

However, using MinHashing to screen for similar pairs has the weakness that the sizes of a read  $R$  and  $P_i$  may vary greatly. Typical databases, like the one used by TaxMapper [1], contain proteins between 10 and 30 000 amino acids (AAs) in length, while one Second Generation Sequencing (SGS) read contains DNA equivalent to about 30 AAs. Consequently, using MinHashing to estimate the Jaccard Similarity of  $R$  and  $P_i$  might result in low resemblance even if  $R$  is perfectly contained in  $P_i$ . This is illustrated in Figure 1.

There are several options to solve this problem. For example Broder [2] proposed to use modulo operations to estimate containment. More recently, Koslicki et al. [3] combined MinHashing with a bloom filter for this purpose. Another established approach is to mitigate the size differences in the first place, by breaking the reference sequences into smaller segments. For example the read mapper VATRAM [4] splits the reference sequence into overlapping windows (of static size).

## Winnowed MinHashing

To mitigate the problems introduced by breaking the references into static sized windows, we developed a variable length indexing approach, based on the winnowing technique [6]. By sliding a window of size  $w$  through all hash values the  $q$ -gram sequence of  $P_i$  and observing the smallest hash value in the window, we obtain segments of the reference which share the same minimizer. This is illustrated in Figure 2.

Consequently, the reference sequence can be partitioned into segments of 1 to  $w$  window starting positions (corresponding to between  $w$  and  $2w - 1$   $q$ -grams). Due to this construction, there is no length  $l \leq w$  subsequence of  $T$ , that has a minimizer that is

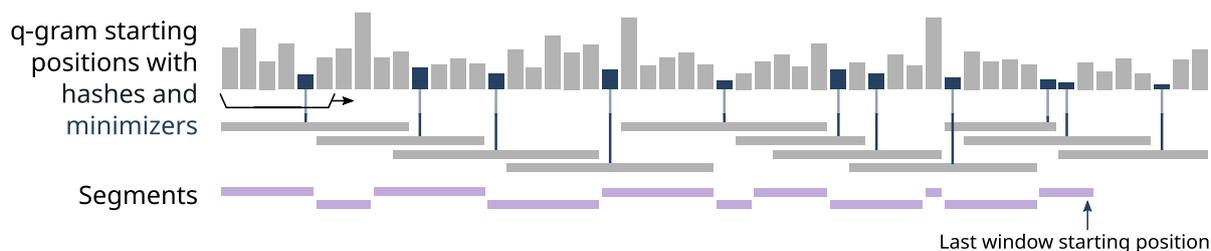


Figure 2: Winnowed segmentation of a reference sequence. Vertical gray bars denote hash values of  $q$ -grams. The dark blue hashes were selected as minimizer using a sliding window of size  $w = 6$ . The horizontal gray bars denote  $q$ -grams dominated by the minimizers. The horizontal purple bars show segments of window start positions that share the same minimizer.

not also the minimizer of a segment. This reduces the size difference problem stated above to a factor of  $\leq \frac{|R|}{2w}$ . However, the exact factor depends on the length of the segments.

## Segment Length Distribution

To judge how well the technique described above mitigates the size difference problem, the distribution of segment lengths needs to be known. We empirically computed the length profiles for the protein database used by TaxMapper, as well as on random DNA sequences. Both did show the characteristic shape illustrated by the gray bars in Figure 3. From a local maximum at position 1, longer segments are continually more unlikely up to a length of 50, which contains the majority of segments. The peak at position 50 corresponds with the window length  $w$  and is caused by a small minimum entering the window and dominating it for the whole window length until it is pushed out again. Segment lengths larger than  $w$  are only possible if the same small minimizer is repeated. This is unlikely for random sequences, but occurs in genomes, due to repetitive sequences.

One of our bachelor students approximated this shape using an exponential function (see [5]). Additionally, we were able to devise a recursive formula to describe the shape that depends on the size of the hash universe  $m$ ,  $q$ -gram size, window size  $w$ , and a cutoff parameter  $k$ . The result of the prediction for several hash universe sizes is illustrated in Figure 3. Even for small hash universes ( $q \approx 500$ ), the predicted distribution closely resembles the empirical distributions.

However, computing the prediction for larger hash universes is very resource intensive and still offers much room for optimization.

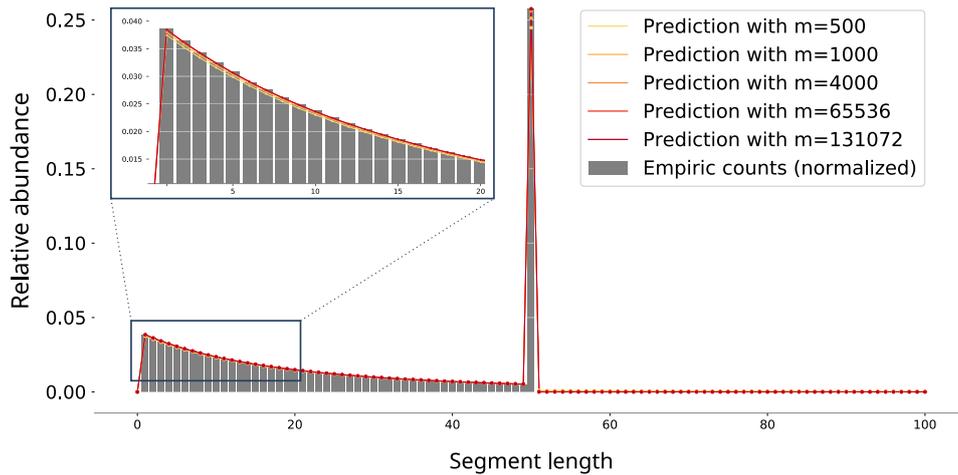


Figure 3: Distribution of segment lengths with  $q = 31$  and  $w = 50$ . Counts from a simulated DNA sequence with  $10^8$  bases are shown as gray bars. The predicted distribution for varying sizes of hash universes  $m$  are shown as yellow and red dots (lines were added to better illustrate the difference between different values of  $m$ ). In the top left corner, a detailed view on the low segment lengths is shown.

## References

- [1] Daniela Beisser, Nadine Graupner, Lars Grossmann, Henning Timm, Jens Boenigk, and Sven Rahmann. Taxmapper: an analysis tool, reference database and workflow for metatranscriptome analysis of eukaryotic microorganisms. *BMC genomics*, 18(1):787, 2017.
- [2] Andrei Z Broder. On the resemblance and containment of documents. In *Compression and Complexity of Sequences 1997. Proceedings*, pages 21–29. IEEE, 1997.
- [3] David Koslicki and Hooman Zabeti. Improving min hash via the containment index with applications to metagenomic analysis. *bioRxiv*, page 184150, 2017.
- [4] Jens Quedenfeld and Sven Rahmann. Variant tolerant read mapping using min-hashing. *arXiv preprint arXiv:1702.01703*, 2017.
- [5] Dennis Rieks. *Eigenschaften von Min-Hashing auf  $q$ -Grammen*. Bachelor’s thesis, TU Dortmund University, 2018.
- [6] Saul Schleimer, Daniel S Wilkerson, and Alex Aiken. Winnowing: local algorithms for document fingerprinting. In *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, pages 76–85. ACM, 2003.





## Subproject C3

Multi-level statistical analysis of  
high-frequency spatio-temporal process data

Katharina Morik      Wolfgang Rhode  
Tim Ruhe

# Bayesian Spectral Analysis for Cherenkov Telescopes

Kai Brügge

Lehrstuhl für Astroteilchenphysik

Technische Universität Dortmund

kai.bruegge@tu-dortmund.de

This report gives a quick explanation on how to calculate energy spectra from data measured by Cherenkov telescopes. I use publicly available data observed by the H.E.S.S, Veritas, MAGIC, and FACT telescopes to fit the spectral energy distribution of the Crab nebula. I use Markov Chain Monte Carlo methods (MCMC) to optimize the likelihoods and quantify the statistical errors.

## 1 IACT Observations

Imaging atmospheric Cherenkov telescopes measure very high energy gamma-rays from distant sources in the universe. The four prominent, and currently operating, telescopes are H.E.S.S, Veritas, MAGIC, and FACT. In a joint effort started in 2016, each collaboration has agreed to publish parts of their data in a common format. [3] The published data includes matrices and tables to describe the instrument response functions as well as measurements of gamma-rays recorded from the Crab Nebula. The actual measurements are encoded as a lists of likely gamma-ray candidates, their estimated energy, and point of origin on the sky. The gamma-ray candidates are divided into signal and background regions using the reflected regions method. [1, 2]

## 2 Statistical Modeling for SED analysis

Spectral energy distributions (SEDs) are of particular interest to many astronomers. SEDs show the energy resolved output of a source over the observed energy range. These curves can often be used to validate or invalidate models of cosmic ray acceleration mechanisms.

Most established analysis of IACT data follow the arguments by Li and Ma [6] to build a statistical model of the measurement process. [5] The likelihood for the expected signal counts per energy bin can be expressed using using a joint poisson for the signal and background region.

$$\mathcal{L} \sim \prod_i \text{Poisson}(\mu_{s,i} + \alpha\mu_{b,i} \mid n_{on,i}) \cdot \text{Poisson}(\mu_{b,i} \mid n_{off,i})$$

Here  $\alpha$  is the ratio between exposures in the signal and background regions,  $\mu_s$  and  $\mu_b$  the expected counts in the signal and background region respectively, and the measured counts  $n_{on}$  and  $n_{off}$ . The  $\mu_s$  are calculated by forward folding a spectral model through the detector response function. Essentially reducing the number of free parameters. Often no such instrument response can be supplied for the background counts as it would require expensive simulations of hadronic interactions in the atmosphere for different observation conditions. This leaves one free parameter per energy bin which is undesirable for minimization and is circumvented by building the profile likelihood. The  $\mu_b$  are disregarded by building the profile likelihood and expressing the  $\mu_b$  as a function of the other parameters by setting

$$\frac{\partial \mathcal{L}}{\partial \mu_b} = 0. \quad (1)$$

The resulting profile likelihood has only as many parameters left as the assumed spectral model. For the Crab Nebula In this analysis I assume a log-parabola spectral form

$$\frac{d\phi}{dE} = \phi_0 \left( \frac{E}{E_0} \right)^{-\Gamma - \beta \log_{10} \left( \frac{E}{E_0} \right)}, \quad (2)$$

since it is a popular choice to approximate the energy spectrum of the Crab Nebula. To sample the posterior prior distributions for the parameters have to be defined. For this analysis I chose the following priors for all datasets

$$\begin{aligned} \Phi_0 &\sim \mathcal{N}_T(\mu = 4, \sigma = 1) \\ \alpha &\sim \mathcal{N}_T(\mu = 2.5, \sigma = 1) \\ \beta &\sim \mathcal{N}_T(\mu = 4, \sigma = 1). \end{aligned}$$

where  $\mathcal{N}_T$  is a normal distribution which is truncated at zero. The high-level statistical model can easily be expressed using probabilistic programming frameworks like PyMC [8]. The forward fold step has been performed using numerical integration, since there is no analytically closed form for the integral of 2. The model could easily be sampled using stochastic samplers like Metropolis. However the number of samples needed to converge is excessively large. I use Hamilton Monte Carlo samplers instead. The difficulty lies in expressing the forward fold step while keeping information about the gradient. I implemented the forward fold step as a computational graph within the Theano framework [9] which allows for gradient based sampling methods like NUTS [4] in conjunction with PyMC.

### 3 Results and Future Work

Figure 1 shows the result of 10 000 samples gathered from the profile likelihood for the MAGIC and the FACT data. The results are compatible with the ones gained using a classical approach to likelihood minimization and error estimation in [7] (to be published in 2019). The profile likelihood, with only 3 free parameters, could have easily been sampled using non-hamilton Monte Carlo approaches. The next step however is to sample the full, non profile, likelihood including the nuisance parameter  $\mu_b$  for each energy bin. This will increase the number of parameters to approximately 50. At the moment the integral of the spectral model (2) is calculated in a numerical way. An analytical solution, or even approximation, will reduce runtime greatly.

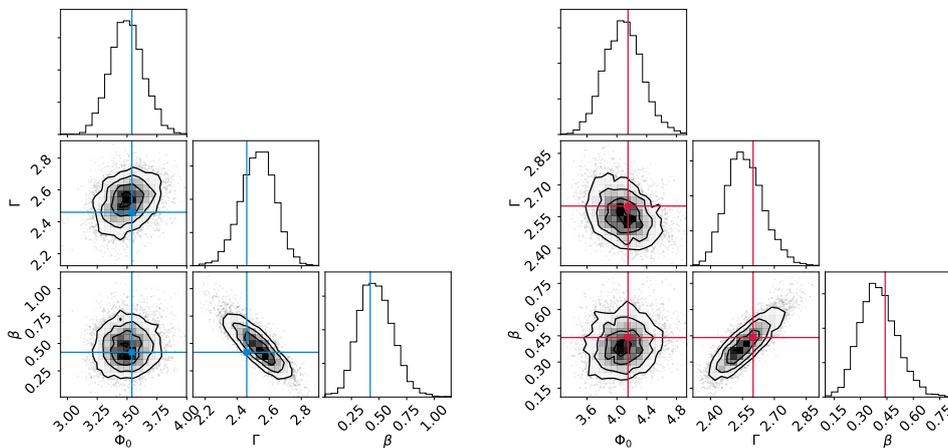


Figure 1: These distributions plots show the samples in the Monte Carlo chain. The blue and red lines indicate the values calculated in [7] using the same data for FACT and MAGIC respectively. The sampled distributions apparently match the maximum likelihood values from [7] closely.

## References

- [1] F. Aharonian et al. Evidence for TeV gamma ray emission from Cassiopeia A. *AAP*, pages 112–120, 2001.
- [2] Berge, D., Funk, S., and Hinton, J. Background modelling in very-high-energy gamma-ray astronomy. *A&A*, 466(3):1219–1229, 2007.
- [3] C. Deil, C. Boisson, K. Kosack, J. Perkins, J. King, P. Eger, et al. Open high-level data formats and software for gamma-ray astronomy. *AIP Conference Proceedings*, 1792(1):id.070006, 2017.
- [4] Matthew Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [5] Johannes King. Hochenergetische gammastrahlung aus dem zentrum der galaxie, 2018.
- [6] T.-P. Li and Y.-Q. Ma. Analysis methods for results in gamma-ray astronomy. *Astroparticle Journal*, 272:317–324, September 1983.
- [7] C Nigro, C Deil, R. Zanin, T. Hassan, J. King, J.E. Ruiz, R. Terrier, K. Brügge, M. Nöthe, R. Bird, T. T. Y Lin, J. Aleksić, C. Boisson, J. L. Contreras, A. Donath, L. Jouvin, N. Kelley-Houskin, B. Khelifi, K. Kosack, and J. Rico. Towards open and reproducible multi-instrument analysis in gamma-ray astronomy. 2018.
- [8] John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. Probabilistic programming in python using pymc3. *PeerJ Computer Science*, 2:e55, April 2016.
- [9] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.

# Deconvolution in Cherenkov Astronomy

Mirko Bunse

Lehrstuhl für Künstliche Intelligenz

Technische Universität Dortmund

mirko.bunse@tu-dortmund.de

Obtaining the probability density function (pdf) of a physical quantity is one frequent objective in experimental physics. In Cherenkov astronomy, one such quantity is the energy of cosmic gamma radiation, which is used to reason about the characteristics of celestial objects. If, however, the relevant quantity cannot be assessed experimentally, it has to be reconstructed from correlated quantities which are measured instead. *Deconvolution* is the reconstruction of the pdf of an in-accessible quantity from secondary, measured quantities. In the last year, we developed a novel unified view on deconvolution algorithms, rephrasing them in the language of data science. Based on our unified formulation, we found a novel stopping condition which guarantees fast convergence.

## 1 Introduction

An accurate and reliable estimate of the sought-after probability density function, e.g. the energy spectrum of an astrophysical particle source, is crucial to understand the underlying physical principles of observed phenomena. In cases where the density of the relevant quantity cannot be accessed experimentally, *deconvolution* (also known as “un-folding”) is applied to estimate this density from correlated quantities that are measured, instead. This estimation is usually aggravated by several deficiencies of the measurement apparatus [1]. Its name stems from the measured pdf  $g : \mathcal{X} \rightarrow \mathbb{R}$  being modeled as a convolution of the relevant pdf  $f : \mathcal{Y} \rightarrow \mathbb{R}$  with a detector response function  $R : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ . To obtain  $f$ , this model of  $g$  has to be inverted—it has to be “de-convolved”.

$$g(x) = \int_{\mathcal{Y}} R(x|y) \cdot f(y) dy \quad (1)$$

## 2 Classical Algorithms

Classical deconvolution algorithms proposed in experimental physics solve a discrete variant  $\mathbf{g} = \mathbf{R}\mathbf{f}$  of the problem statement from Eq. 1. Estimates  $\hat{\mathbf{f}}$  obtained this way tend to exhibit a large variance, due to the usually poor condition of the matrix  $\mathbf{R}$ . We identified the discretization of  $\mathbf{g}$  to be a clustering problem. Supervision of the clustering can improve the critical matrix condition and lead to more stable results [2].

Our framework [4] unifies the most well-known deconvolution algorithms RUN [1], IBU [5], and, as a work in progress, an SVD-based approach [6]. Thus, it discloses similarities and differences between algorithms, guiding practitioners in algorithm selection and paving the way for new combinations of algorithmic aspects. Even though the presented algorithms are widely adopted in the physics community, their performances have not been compared with each other previously. Our work closes this gap but also reveals that all of these algorithms can produce similarly accurate estimates in the domain of Cherenkov astronomy, if their meta-parameters are tuned well [4].

## 3 Deconvolution through Supervised Learning

The *Dortmund Spectrum Estimation Algorithm* (DSEA) [7] takes another path than the classical algorithms by re-phrasing deconvolution as a supervised learning task. Here, the estimate is aggregated from classifier predictions on the relevant quantity of individual examples. It is then improved by iterating the reconstruction, updating the density of the training set—which is used to obtain the classifier—with the latest estimate.

DSEA underlies the intriguing idea that  $\mathbb{P}(Y \equiv i)$  can be recovered from a classifier's confidence. In addition, a uniform prior is imposed on  $\mathbf{x}$ , i.e.  $\hat{\mathbb{P}}(\mathbf{X} = \mathbf{x}) = \frac{1}{N}$ . The outcome is the DSEA estimator, which estimates  $\mathbf{f}$  from confidence values.

$$\hat{\mathbf{f}}_i = \frac{1}{N} \sum_{n=1}^N c_{\mathcal{M}}(i | \mathbf{x}_n) \quad (2)$$

The advantages of DSEA are manifold: Un-discretized inputs  $\mathbf{X}$  are supported, so that the clustering required for classical algorithms can be circumvented. Moreover, aggregating contributions (confidences) of individual examples will enable us to study the deconvolution result as a function of other variables, e.g. zenith angle and time, in future work. Last but not least, the algorithm paves the way for a tightened collaboration between physicists and machine learning experts. However, the original DSEA has been shown to diverge from the optimal solution after having found a suitable estimate [4].

## 4 Adaptive Steps for Fast Convergence

To cope with the divergence of DSEA and to reduce the number of iterations needed, we proposed the extended algorithm DSEA<sup>+</sup> [4], which embraces three aspects:

- Instead of weighting each example with the estimated probability of its discrete state, each example is weighted with the *ratio* between this estimated probability and the corresponding probability in the un-weighted training set. This modification stops the divergence almost entirely.
- The steps  $\mathbf{p}^{(k)} = \hat{\mathbf{f}}^{(k)} - \hat{\mathbf{f}}^{(k-1)}$  taken by DSEA are scaled with some  $\alpha^{(k)} \geq 0$ , which leads to the update rule  $\hat{\mathbf{f}}^{(k)+} = \hat{\mathbf{f}}^{(k-1)} + \alpha^{(k)} \cdot \mathbf{p}^{(k)}$ . At first, we investigated some simple strategies to choose  $\alpha^{(k)}$ , namely constant and decaying step sizes [3]. Even though these strategies enforce convergence, they are hard to tune and they slow down the algorithm drastically.
- Therefore, a near-optimal step size  $\alpha_{\text{RUN}}^{(k)}$  is chosen, based on the regularized objective function  $\ell_r$  from RUN. Unlike in RUN,  $\ell_r$  is minimized only in the search direction that is determined by DSEA. We observe that this adaptive strategy converges as soon as the likelihood can not be improved in the given search direction. Moreover, the algorithm makes use of step sizes greater than one, thus approaching the optimal solution faster than the original DSEA.

$$\alpha_{\text{RUN}}^{(k)} = \arg \min_{\alpha \geq 0} \ell_r \left( \hat{\mathbf{f}}^{(k-1)} + \alpha \cdot \mathbf{p}^{(k)} \right) \quad (3)$$

## 5 Future Work

We will employ DSEA<sup>+</sup> in physical analyses of real-world telescope observations, also leveraging that this algorithm is the first to enable time-dependent (or angle-dependent) deconvolution. Moreover, we seek to re-phrase deconvolution once again—this time not in terms of classification but regression. The currently taken classification approach has been motivated by the availability of confidence values expressing the uncertainty of predictions. Lifting DSEA<sup>+</sup> to regression will in turn require an accurate assessment of uncertainty in regression predictions. If successful, this new approach will produce continuous estimates of the sought-after probability density, which are more versatile than the discrete estimates produced by the existing classification-based approach.

Concerning the other plans of the C3 project—simulation, distributed learning, and tailored deep learning—deconvolution is always a key aspect: In almost any analysis of real-world data, deconvolution is the final step. It summarizes to the practitioner what the telescope has seen, bridging from data analysis to physical interpretation. Therefore, we must evaluate our advances in other areas also in terms of deconvolution accuracy.

Our deconvolution software is starting to be adopted by physicists. It is available as an open source package for Julia and Python, with additional information hosted at <https://sfb876.tu-dortmund.de/deconvolution>.

## References

- [1] V. Blobel. An unfolding method for high energy physics experiments. In *Adv. Stat. Techniques in Part. Phys.*, pages 258–267, 2002.
- [2] M. Börner, T. Hoinka, M. Meier, T. Menne, W. Rhode, and K. Morik. Measurement/simulation mismatches and multivariate data discretization in the machine learning era. In *Proc. of the ADASS XXVII*, 2017. In press.
- [3] Mirko Bunse, Nico Piatkowski, and Katharina Morik. Towards a unifying view on deconvolution in Cherenkov astronomy. In Rainer Gemulla, Simone Paolo Ponzetto, Christian Bizer, Margret Keuper, and Heiner Stuckenschmidt, editors, *Lernen, Wissen, Daten, Analysen (LWDA) conference proceedings*, pages 73–77, 2018.
- [4] Mirko Bunse, Nico Piatkowski, Tim Ruhe, Wolfgang Rhode, and Katharina Morik. Unification of deconvolution algorithms for Cherenkov astronomy. In *Proc. of the 5th DSAA*, pages 21–30. IEEE, 2018. In press.
- [5] G. D’Agostini. Improved iterative Bayesian unfolding. *arXiv:1010.0632*, 2010.
- [6] A. Hoecker and V. Kartvelishvili. SVD approach to data unfolding. *Nucl. Instrum. Methods Phys. Res. A*, 372(3):469–481, 1996.
- [7] T. Ruhe, T. Voigt, M. Wornowizki, M. Börner, W. Rhode, and K. Morik. Mining for spectra – the Dortmund spectrum estimation algorithm. In *Proc. of the ADASS XXVI*, 2016. In press.

# Cascade Reconstruction in IceCube using Convolutional and Generative Neural Networks

Mirco Hünnefeld  
Lehrstuhl für Experimentelle Physik E5b  
Technische Universität Dortmund  
mirco.huennefeld@tu-dortmund.de

Reliable and accurate reconstruction methods are vital to the success of high-energy physics experiments such as IceCube. Machine learning based techniques, in particular deep neural networks, can provide a viable alternative to maximum-likelihood methods. Most common neural network architectures originate from non-physical domains such as image recognition. While these methods can enhance the reconstruction performance in IceCube, there is much potential for tailored techniques. In the typical physics use-case, many symmetries, invariances and prior knowledge exist in the data, which are yet to be exploited by current network architectures. Novel and specialized deep learning based reconstruction techniques are desired which can leverage the physics potential of experiments like IceCube. A new approach using generative neural networks for the reconstruction of cascade-like events in IceCube is presented.

## 1 Event Reconstruction in IceCube

A key challenge to the success of experiments such as IceCube is the reliable and accurate reconstruction of events. In IceCube, further challenges arise as the detector is situated at the geographic South Pole where resources are limited. This results in a dilemma since performance is often paired with computational complexity. But even for offline reconstructions, the computational complexity of the most advanced maximum-likelihood methods can render these intractable and hence limit the physics potential.

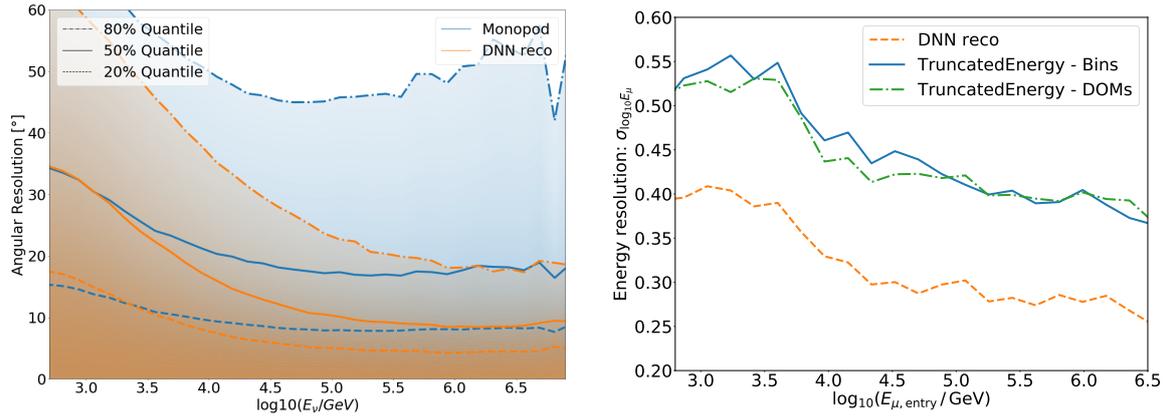


Figure 1: On the left, the angular resolution for cascade events is shown for the current state-of-the-art reconstruction method (Monopod [3]) and for the newly developed method based on deep neural networks (DNN reco). On the right, the muon energy resolution is shown for the current standard methods [3, 4] and the deep learning based method.

Machine learning based methods, in particular deep neural networks, might help to overcome these obstacles. The use of convolutional neural networks (CNNs) [1] can greatly enhance the reconstruction performance in IceCube as shown in [2]. Despite their success, convolutional architectures have considerable limitations. They assume translational invariance and the data to be aligned on a regular grid. These assumptions are only approximately fulfilled in IceCube. Most importantly though, prior knowledge cannot easily be exploited by the network architecture as it could, for instance, in maximum-likelihood methods. Novel deep learning based methods specifically tailored to the needs in high-energy physics experiments such as IceCube are needed.

## 2 Convolutional and Generative Neural Networks

The previously developed reconstruction method based on CNNs [2] was extended to incorporate classification and regression tasks for all event topologies in IceCube. In addition, extensive tests were performed to validate the performance and to quantify the effect of systematic uncertainties in the Monte Carlo simulation. As of writing this, the CNN approach provides the best muon energy and cascade directional reconstruction in IceCube as illustrated in fig. 1. It can improve the resolution by up to 50 % while reducing the runtime by two to three orders of magnitude. The developed method has already been used in a search for neutrino emission from the Galactic Plane where it improved the sensitivity by almost a factor of 2. Details will be provided in a future publication.

First steps have been taken to incorporate prior knowledge into the network architecture

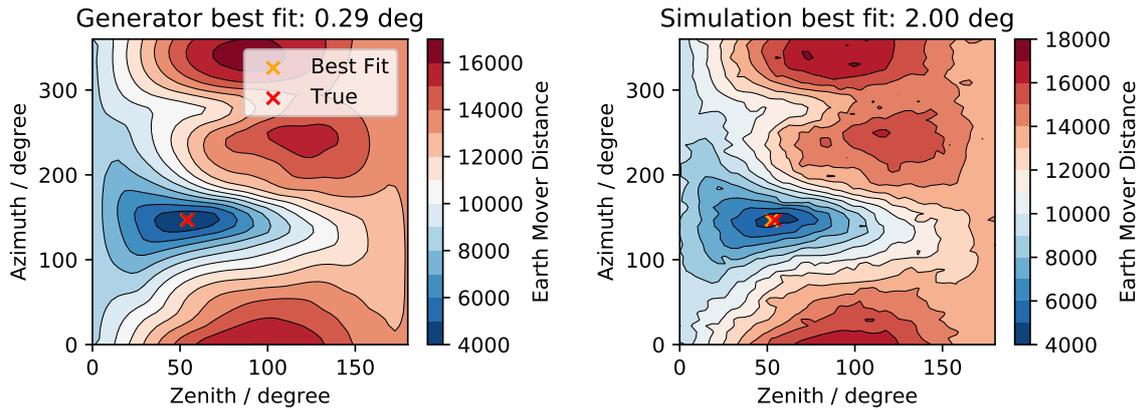


Figure 2: The parameter space is scanned to obtain the best fit value for the cascade direction. In contrast to the simulation (right), the scan over the generator output (left) is less noisy and provides smooth gradients for a minimization.

as it is done for maximum-likelihood-methods [5]. For a maximum-likelihood-method, one generally compares measured data to some expectation value. This can for instance be the number of measured photons at a specific photomultiplier. The expected number of photons is then obtained from Monte Carlo simulations. This process, while being the most accurate, is extremely slow and not always tractable. Reconstruction of a single event may take hours to days in an approach like this.

Instead of performing computing intensive simulations, a generative network can be used to approximate these. A neural network, the generator, is trained to predict the expected waveforms at each DOM for a given cascade hypothesis. Once the generator network is trained, it can then be used in reverse mode. The waveforms for a given event can now be compared to the generated waveforms for a given hypothesis. A distance measure is computed between the generated and true waveforms. This distance (negative log likelihood,  $\chi^2$ , earth mover distance, etc.) is a function of the cascade hypothesis and can be minimized analogous to a maximum-likelihood approach.

An example for an event with two free parameters, the zenith and azimuth direction of the cascade, is shown in fig. 2. The distance measure, in this case the earth mover distance, is calculated for every (zenith, azimuth) pair in order to obtain a scan of the landscape. The landscape of the generator network is much better behaved and provides smooth gradients, which is not the case for the simulated events. For this particular event, the reconstruction via the generator network provides a reconstructed direction which is only  $0.29^\circ$  away from the true cascade direction as opposed to the simulation which only achieves a resolution of  $2^\circ$ . The simulation required many CPU and GPU hours, while the generator network (once trained) only required  $\mathcal{O}(s)$ . In contrast to the simulation, the generator network is fully differentiable, hence, gradient descent can be used to maximize the likelihood. In this case, the per event reconstruction time reduces

to  $\mathcal{O}(100 \text{ ms})$ .

In this approach, translational invariance in the cascade hypothesis can be exploited rather than the approximative translational invariance in measured data which is affected by the detector acceptance and ice inhomogeneities. Moreover, irregularities in the detector grid are now naturally accounted for and prior knowledge such as how the waveforms are parameterized can easily be included in the generator architecture to reduce the number of free parameters. Challenges remain once the dimensionality is increased to the full cascade hypothesis consisting of the vertex position and time, the initial cascade direction and energy. Nevertheless, the results so far are very promising.

### 3 Conclusion and Outlook

Standard neural network architectures such as CNNs are well suited for the application in IceCube and can further improve the event reconstruction accuracy, while greatly reducing the per event runtime. Despite their success, these methods have considerable limitations. In the physics use-case, extensive knowledge typically exists about the data generation process, constraints, and physics laws the data obeys. Yet, current deep learning architectures can not fully exploit this information. Novel and tailored deep learning based reconstruction techniques are therefore desired which can combine the strengths of existing architectures and maximum-likelihood methods. The presented approach using generative neural networks provides a promising step towards these goals.

### References

- [1] Y. LeCun et al., *Handwritten Digit Recognition with a Back-Propagation Network*, Adv Neural Inf Process Syst **2**, 396–404 (1990)
- [2] M. Huenefeld (IceCube Collaboration), *Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube*, PoS **ICRC2017**, 1057 (2017)
- [3] M. Aartsen et al. (IceCube Collaboration), *Energy reconstruction methods in the IceCube neutrino telescope*, JINST **9**, P03009 (2014)
- [4] R. Abbasi et al. (IceCube Collaboration), *An improved method for measuring muon energy using the truncated mean of  $dE/dx$* , Nucl. Instrum. Meth. **A703**, 190-198 (2013)
- [5] A. Flores and M. Huenefeld, *RISE Germany Internship: Generative Neural Network Reconstruction of IceCube Cascade-like Events*, SFB876-C3, TU Dortmund, Report No. **5**, (2018)

# Search for High-Energy Astrophysical Tau-Neutrinos in Seven Years IceCube Data

Maximilian Meier  
Lehrstuhl für Experimentelle Physik V  
Technische Universität Dortmund  
maximilian.meier@tu-dortmund.de

In this work a potential astrophysical flux of tau neutrinos is investigated. The measurement of tau neutrinos would be a clear astrophysical signal since no tau neutrino flux from the atmosphere is expected. Tau neutrinos can be detected with the IceCube detector by their unique double-cascade signature at high-energies. But the signal is buried in a large amount of background events with a signal-to-background ratio of about  $\mathcal{O}(1 : 10^{10})$  at trigger level. To remove these background events methods from the field of machine learning are applied. A detection of tau neutrinos would require a very pure sample to achieve a high significance and a very efficient analysis to obtain sufficient statistics in limited time of measurement. This report presents the results achieved in the year 2018.

## 1 Introduction

A physical motivation, a detailed introduction to the IceCube detector and a description of all event signatures relevant for this analysis can be found in [7].

To analyse the astrophysical flux of tau neutrinos a sample with a high expected fraction of tau neutrinos has to be obtained at first. This can be achieved by a two-step event selection. The first step is the selection of events with a double pulse signature in at least one of IceCubes Digital Optical Modules (DOMs), which is described in [7] and [6]. The second step is the selection of cascade-like events over a background of track-like events

originating from charged current muon neutrino interactions and atmospheric muons. The process of feature selection for this classification task is illustrated in [6].

## 2 Cascade Selection

The starting point for the cascade selection are all events that deposit at least 2000 photoelectrons, corresponding to deposited energies of roughly 10 TeV, and that show a double pulse waveform in at least one DOM throughout the event. These requirements reduce backgrounds from single cascades (charged current electron neutrino events and all flavor neutral current events) to a subdominant level with respect to the sought after signal, charged current tau neutrinos. The signal is still dominated by track-like backgrounds, that can also produce double pulses, but fortunately show different topologies in the detector.

After applying the process described in [6] to select suitable features for this classification task a Random Forest [3] is trained to solve this task. The signal component is splitted up before training in contained and uncontained event and only those where the tau neutrino interaction vertex and the tau decay vertex are both within the detector boundaries are considered as signal in the training. Fig. 1 shows the result of the classification in form of the expected event rate for each component as a function of the Random Forest classification score.

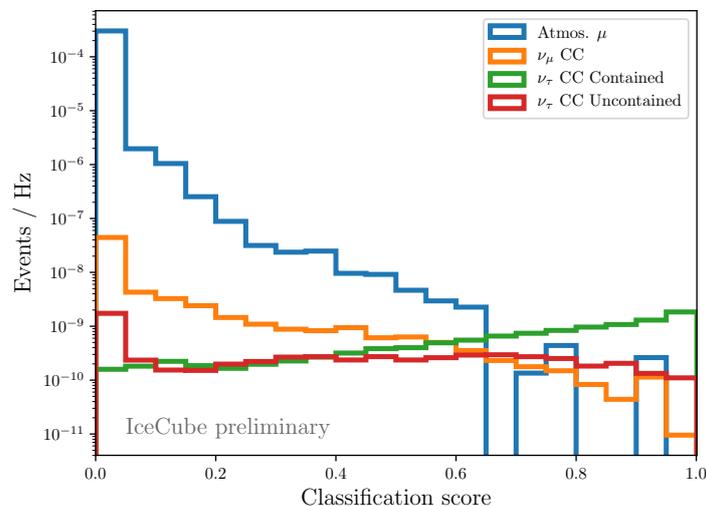


Figure 1: Expected event rates for different simulated components as a function of the Random Forest classification score.

The classification score cut to obtain the final sample is optimized via the model rejection factor [5],

$$\text{MRF} = \frac{1}{n_s} \sum_{n_{\text{obs}}=0}^{\infty} \mu_{90}(n_{\text{obs}}, n_b) \frac{n_b^{n_{\text{obs}}}}{n_{\text{obs}}!} \exp(-n_b), \quad (1)$$

where  $\mu_{90}(n, n_b)$  is the 90% poisson upper limit for  $n$  observed events with a background expectation of  $n_b$  and zero expected signal. This optimisation yields a classification score cut of 0.62 and a MRF of  $\sim 1.6$  for an astrophysical flux assumption [1] of

$$\Phi_{\nu}(E) = 0.90 \cdot 10^{-18} \cdot \left( \frac{E}{100 \text{ TeV}} \right)^{-2.13} \frac{1}{\text{GeV cm}^2 \text{ sr s}}. \quad (2)$$

The model rejection factor is directly connected to the average upper limit that can be set on the flux assumption it is calculated for

$$\bar{\Phi}_{\nu,90}(E) = \Phi_{\nu}(E) \cdot \text{MRF} = \Phi_{\nu}(E) \cdot \frac{\bar{\mu}_{90}(n_{\text{obs}}, n_b)}{n_s}. \quad (3)$$

This results in a signal expectation of roughly two events in 7 years of livetime and a background expectation of slightly less than one event. This is an increase in signal rate compared to a previous double pulse analysis conducted by the IceCube collaboration [2] by a factor of 2.14, while keeping a similar background rate (when comparing event rates for similar astrophysical flux assumptions).

### 3 Analysis Sensitivity

The tau neutrino flux normalisation will be measured with a Poisson likelihood fit, where the likelihood is given by

$$\mathcal{L}(n|\lambda) = \prod_{i=0}^N \mathcal{P}_{P_{B,i} + \lambda P_{S,i}}(n_i), \quad (4)$$

with the number of measured events  $n$  as well as the binned PDFs for signal and background  $P_S$  and  $P_B$ . Since the expected amount of total events is very low, the background is kept fixed and only the signal normalisation is varied with a fixed spectral index. To estimate the sensitivity of the analysis pseudo-experiments assuming no true signal are performed. The sensitivity is defined as the average upper limit of these pseudo-experiments. The confidence intervals presented are constructed as unified confidence intervals with the Feldman-Cousins method [4]. Fig. 2 shows the distribution of 90% upper limits for 1000 pseudo-experiments assuming the astrophysical flux described in eq. 2. The resulting sensitivity is  $\sim 1.1 \times 10^{-18}/(\text{GeV cm}^2 \text{ sr s})$ .

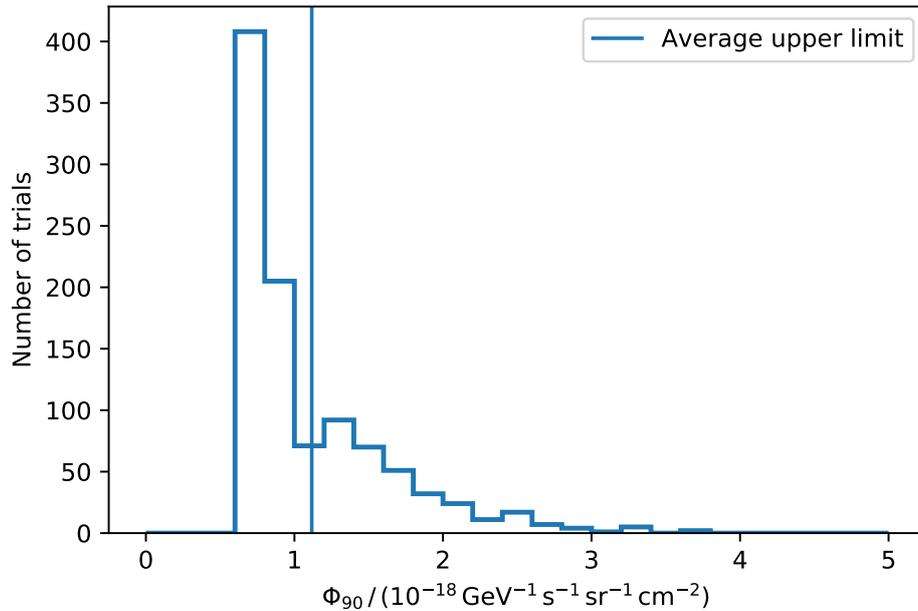


Figure 2: Distribution of 90 % upper limits on the tau neutrino flux normalisation for 1000 background only pseudo-experiments.

## References

- [1] M. G. Aartsen et al. In: *Astrophys. J.* 833.1 (2016), p. 3.
- [2] M. G. Aartsen et al. “Search for Astrophysical Tau Neutrinos in Three Years of IceCube Data”. In: *Physical Review D* 93 (2016), p. 022001.
- [3] Leo Breiman. “Random Forests”. In: *Machine Learning* 45.1 (2001), pp. 5–32.
- [4] Gary J. Feldman and Robert D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev. D* 57 (7 Apr. 1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873. URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.3873>.
- [5] Gary C. Hill and Katherine Rawlins. “Unbiased cut selection for optimal upper limits in neutrino detectors: The Model rejection potential technique”. In: *Astropart. Phys.* 19 (2003), pp. 393–402. DOI: 10.1016/S0927-6505(02)00240-2. arXiv: astro-ph/0209350 [astro-ph].
- [6] M. Meier. *Cascade Selection for a Search for High-Energy Astrophysical Tau-Neutrinos in IceCube Data*. Technischer Report für das SFB 876 Gradiertenkolleg. TU Dortmund University, 2017.
- [7] M. Meier. *Search for high-energy astrophysical Tau-Neutrinos in IceCube Data*. Technischer Report für das SFB 876 Gradiertenkolleg. TU Dortmund University, 2016.

# FACT-Tools 2018 Development Report

Maximilian Nöthe  
Lehrstuhl für Experimentelle Physik 5b  
Technische Universität Dortmund  
maximilian.noethe@tu-dortmund.de

The FACT-Tools, an extension of the `streams` frame-work to analyze the data of the First G-APD Cherenkov Telescope (FACT), saw major updates in 2017 and 2018. Since the last release version `v0.17.2`, released on 5.12.2016, many new features, code cleanups, and bug fixes were implemented. Between `v0.17.2` and the current version `v1.1.2`, released on 17.12.2018, 992 commits by 8 contributors were merged into master. The most important changes include a full system of coordinate transforms from ICRS via horizontal to the FACT camera frame, a new reader for fits files, a new inputstream to read runwise simulation metadata, a single xml steering file for both simulations and observations, a service that provides the correct gain for each run automatically, a service that finds the correct drs file for each run, and a service that finds bright stars in the field of view.

## 1 Introduction

The FACT-Tools are an extension to the `streams` framework [3], developed to analyze the observed as well as simulated data of the First G-APD Cherenkov Telescope (FACT) [1, 2]. FACT is an Imaging Atmospheric Cherenkov Telescope (IACT), observing since October 2011 at the Observatorio del Roque de los Muchachos on the Canary Island of La Palma, Spain. Every night, given good weather conditions, FACT takes between 400 GB and 1 TB of raw data.

To reconstruct the particle type, energy and origin of the recorded air showers, several steps have to be performed.

1. Calibration of the raw time series in each pixel
2. Extracting the number of photons and their mean arrival time in each pixel
3. Removal of pixel not likely to contain Cherenkov signal
4. Parameterization of the resulting images
5. Reconstruction of energy, particle type and origin from the image parameters

Steps 1 through 4 are performed using `FACT-Tools`, resulting in one fits file per observation run, corresponding to one or five minutes of data taking. A python-based machine-learning reconstruction performs step 5.

Since version `v0.17.2` [5], 992 commits were merged into the master branch:

Table 1: Commits into master since `v0.17.2`

Commits	Author	Institute
514	Maximilian Nöthe	Exp. Physik 5b, TU Dortmund
196	Jens Buß	Exp. Physik 5b, TU Dortmund
126	Kai Brügge	Exp. Physik 5b, TU Dortmund
85	Michael Bulinski	Exp. Physik 5b, TU Dortmund
28	Sebastian A. Mueller	IPP, ETH Zürich
23	Dominik Neise	IPP, ETH Zürich
19	Hendrik Hildebrandt	Exp. Physik 5b, TU Dortmund
1	Christian Bockermann	LS8 Informatik, TU Dortmund

## 2 New Features

In this section, the most important new features in `FACT-Tools v1.1.2` [4] compared to `v0.17.2` are described.

**Coordinate Frames (v1.0.0)** A system of coordinate transforms was implemented. Allowing transformations between equatorial, horizontal and FACT camera coordinates. The transformations from equatorial to horizontal and vice versa implement first order precession corrections which makes them accurate to a level below 30 arcseconds, which is approximately 10 times smaller than a FACT camera pixel. An example on how to use this coordinate transforms can be found in Listing 1.

Listing 1: Get the current position of the Crab Nebula in camera coordinates for a given telescope pointing direction.

```
ZonedDateTime time = ZonedDateTime.now();
EarthLocation location = EarthLocation.FACT;
double focalLength = 4889.0;
HorizontalCoordinate pt = HorizontalCoordinate.fromDegrees(5.0, 0.0);

EquatorialCoordinate eq = EquatorialCoordinate.fromDegrees(83.63, 22.01);
HorizontalCoordinate hz = eq.toHorizontal(time, location);
CameraCoordinate camera = hz.toCamera(pt, focalLength);
```

**New FITS reader (v1.0.0)** The new `hdureader` package provides a much improved reading of fits files, including FACT's custom compressed zfits data files. With the fixes included in v1.1.2, every FITS file produced by the FACT should be read correctly by FACT-Tools.

**CeresStream (v1.0.0)** A new stream was implemented to read important runwise simulation metadata from a second inputfile per simulation run.

**GainService (v1.0.0)** A new service was implemented, making it easy to get the correct gain for a data run. Before, FACT-Tools assumed the gain was constant over time and used a static resource file for game calibration. This assumption turned out to be false and was fixed by providing a database file with measured gains for each night.

**Merged files for simulations and observations (v1.0.0)** There are several differences between simulated and observed FACT data, that require different steps to be performed, although the major part of the analysis is exactly the same. Before FACT-Tools v1.0.0, every analysis step had a steering file each for observations and simulations, resulting in many duplications and making it hard to spot differences and to maintain. For v1.0.0, these xmls got merged into a single xml, branching where necessary.

**Adding measured noise to simulated air showers (v1.1.0)** FACT's simulation chain implements on overly simplistic model of the electronics noise and the night sky background photons. To overcome this, a processor to superimpose measured noise on simulated air shower signal was developed, resulting in labeled dataset with more realistic noise behaviour.

**DrsFileService (v1.1.0)** The calibration of FACT's raw data requires an inputfile with calibration constants that should be as close in time to the data file as possible. For batch processing, this file is selected upfront, this is however not possible for realtime analysis, where the calibration coefficients need to be updated several times per night. For this, a new service was implemented, yielding the correct calibration constants for the current run.

**StarService, (v1.1.0)** Bright stars in the field of view induce additional noise in the camera. To ignore pixels that survive the cleaning process only due to this higher noise level, a new service was implemented that yields bright stars from the Yale Bright Stars catalogue currently in the field of view.

### 3 Bug Fixes and Smaller Improvements

In versions v1.1.2, a bug in the hdureader that prevented correct reading of zfits file having ZTITLELEN > 1, was fixed. Before, only the first event in each tile was read and then repeated.

Before v1.0.2, FACT-Tools would return an exit code of 0, even if an exception was thrown.

Before v1.1.0 the Timespread processor would fail if given an empty pixelset.

v1.0.0 also added the possibility to write header keywords into fits files.

All output keys of the standard analysis are now in snake\_case.

### References

- [1] H Anderhub et al. "Design and operation of FACT—the first G-APD Cherenkov telescope". In: *Journal of Instrumentation* 8.06 (2013), P06008.
- [2] Adrian Biland et al. "Calibration and performance of the photon sensor response of FACT—the first G-APD Cherenkov telescope". In: *Journal of Instrumentation* 9.10 (2014), P10012.
- [3] Christian Bockermann and Hendrik Blom. *The streams Framework*. Tech. rep. 5. TU Dortmund University, Dec. 2012.
- [4] Kai Brügge et al. *fact-project/fact-tools: v1.1.2*. Dec. 2018. DOI: 10.5281/zenodo.2386762.
- [5] Maximilian Nöthe et al. *FACT-Tools Version v0.17.2*. Apr. 2018. DOI: 10.5281/zenodo.1212389.

# Jet Kinematic Analysis using VLBI Images

Kevin Schmidt

Experimentelle Physik 5

Technische Universität Dortmund

kevin3.schmidt@tu-dortmund.de

A large number of extragalactic sources are active galactic nuclei. These can be divided into several subclasses depending on the observed characteristics. A central feature for the classification is the viewing angle of the sources' jet relative to our line of sight.

One aim of this analysis is to study the jet kinematics of candidates of TeV radio galaxies using VLBI images. Determined jet speeds are used to calculate constraints on the viewing angle of the candidates. The results help to define the characteristics of the source population and to learn about acceleration mechanisms of active galactic nuclei in general.

New radio interferometers like LOFAR and SKA observe large parts of the sky on short timescales. For the study of the obtained detailed images improved analysis methods have to be developed since existing ones cannot handle a large amount of data. As a first step towards an automatized analysis, this work is developing a pipeline for jet kinematic analysis.

## Radio Interferometry

Radio interferometry enables high-resolution imaging of extragalactic sources on milliarc-second resolution [5]. Thus, it provides information about their innermost structures, their source classes, and their emission mechanisms. Radio interferometric data can be used to determine the viewing angle of active galactic nuclei. The viewing angle is an important feature for the classification of such sources. Therefore, it is of particular interest.

New radio interferometers like the Low-Frequency Array (LOFRA) [6] and the upcoming Square Kilometer Array (SKA) [1] obtain detailed images of large parts of the radio sky. These telescopes have an improved sensitivity compared to previous experiments, which is achieved by a large number of antennas. The resulting increase in data volumes of 100 PB/day in the case of the SKA is a big challenge for data science. Existing analysis software has to be adapted to be able to perform calibration and imaging of the low-level data in reasonable time. It is therefore essential to develop an automated analysis chain. With the help of the high-level data, it is possible to resolve the extended jet structures of active galactic nuclei and analyze their jet characteristics.

## Study of Jet Characteristics of TeV Radio Galaxies

In this work, high-level radio images are analyzed with the aim to estimate the viewing angle of TeV radio galaxy candidates. To gather information about the evolution of radio-loud sources and their jets, the Monitoring Of Jets in Active galactic nuclei with VLBA Experiments (MOJAVE) program measures these sources at 15 GHz on a regular basis [2]. The obtained images are well suited to compute the velocity of the jets and to determine constraints on the viewing angle of the sources. For the automatization of the calculations, a pipeline for jet kinematic analysis was developed in python.

### Pipeline for Jet Kinematic Analysis

The first task of the analysis is to remove remaining PSF structures from the radio images. Therefore, source models are built iteratively to obtain clean images. These are utilized for the model fitting of the jets. Two-dimensional Gaussian components are used to get a simplified representation of the sources' jet morphologies. The parameters of the Gaussians are adjusted with a chi-square fit. Both tasks are performed using DIFMAP [4].

Figure 1 illustrates the dynamic evolution of the TeV radio galaxy 3C 264. Clean images overlaid with Gaussian components are shown for different observation epochs. The spatial space between the epochs corresponds to the period between the observations. Identified components, which are used for jet kinematic calculations, are shown in color. A linear regression is performed to determine their velocities. Another intrinsic jet parameter is the jet to counter jet flux ratio, as it is related to the Doppler factor of the source. Together with the velocities of the components, it can be used to constrain the viewing angle. Figure 2 displays two areas. The blue one is related to the components' velocity. The flux ratio between jet to counter jet limits the orange one. The overlapping area

indicates the allowed parameter space. Thereby, the upper intersection point marks the upper limit for the viewing angle.

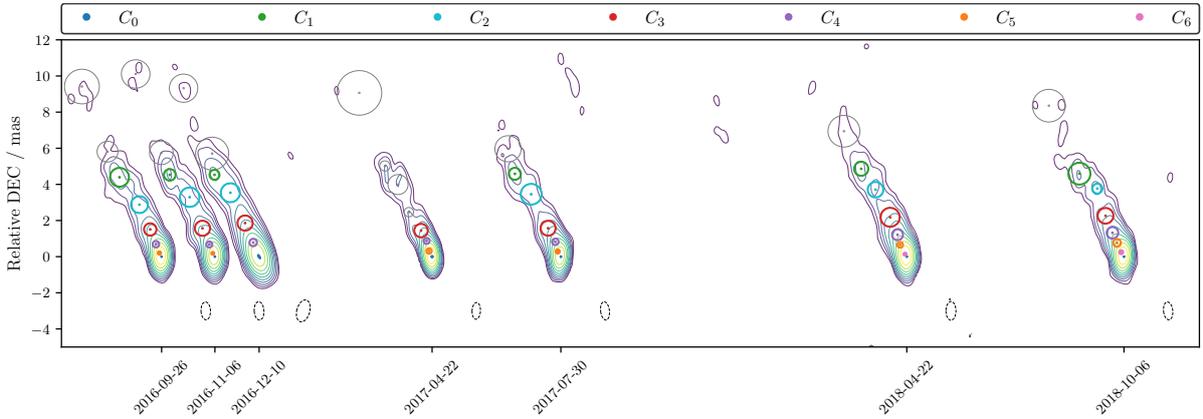


Figure 1: Dynamic evolution of the TeV radio galaxy 3C 264. Clean images are overlaid with model components created in DIFMAP. Identified components are displayed in color. Gray components are not used for calculations. Contours of the clean images begin at  $4\sigma$  and increase logarithmically. The PSF of the individual epochs is represented by black dashed ellipses.

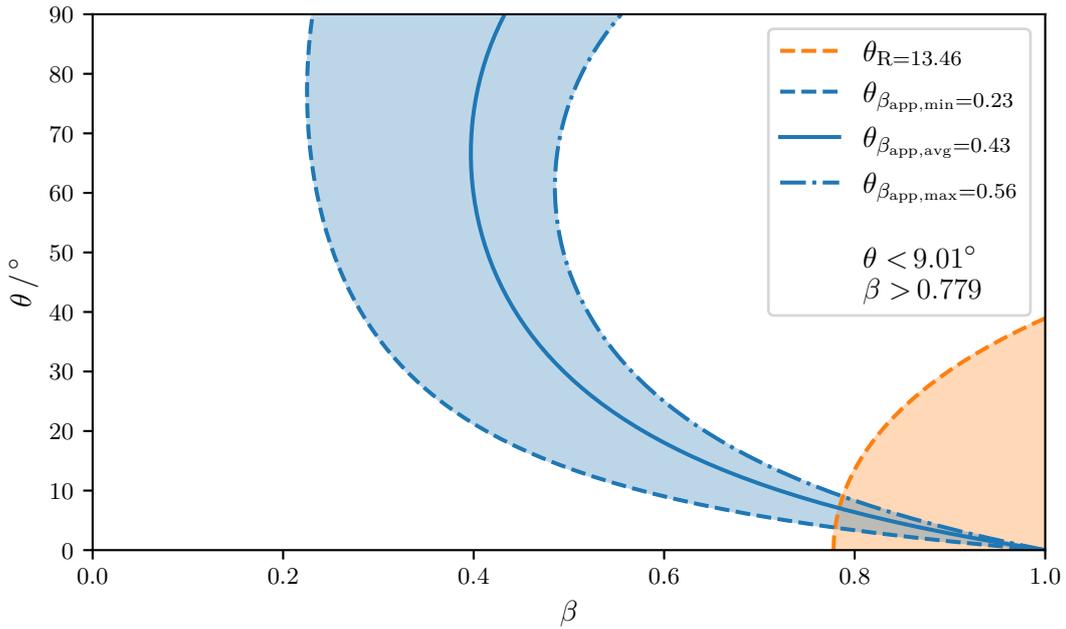


Figure 2: Viewing angle  $\theta$  against jet speed  $\beta$ . The blue surface is limited by the speed of the components. The flux ratio between jet to counter jet limits the orange surface. The top intersection point of the two surfaces gives an upper limit for the viewing angle  $\theta$ , the lower intersection point marks a lower limit for the jet speed  $\beta$ .

All required calculations are combined in one pipeline making manual intervention hardly necessary.

## Outlook

As mentioned in the introduction, new radio interferometers like LOFAR and SKA produce a large number of images, which cannot be analyzed manually. Therefore, it is essential to develop automated pipelines for the modeling and classification of VLBI observations. These will help to increase populations of rare source types and allow to improve our general understanding of astrophysical acceleration mechanisms. Testing recently developed software like WSCLEAN [3] and machine learning approaches using neural networks is currently work in progress.

## References

- [1] P. E. Dewdney et al. The square kilometre array. *Proceedings of the IEEE*, 97(8):1482–1496, Aug 2009.
- [2] M. L. Lister et al. MOJAVE. XV. VLBA 15 GHz Total Intensity and Polarization Maps of 437 Parsec-scale AGN Jets from 1996 to 2017. *The Astrophysical Journal Supplement Series*, 234:12, January 2018.
- [3] A. R. Offringa et al. wsclean: an implementation of a fast, generic wide-field imager for radio astronomy. *Monthly Notices of the Royal Astronomical Society*, 444(1):606–619, 2014.
- [4] M. C. Shepherd, T. J. Pearson, and G. B. Taylor. DIFMAP: an interactive program for synthesis imaging. In *Bulletin of the American Astronomical Society*, volume 26, pages 987–989, May 1994.
- [5] G. B. Taylor, C. L. Carilli, and R. A. Perley, editors. *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, 1999.
- [6] van Haarlem, M. P. et al. Lofar: The low-frequency array. *A&A*, 556:A2, 2013.

This research has made use of data from the MOJAVE database that is maintained by the MOJAVE team (Lister et al., 2018, ApJS, 232, 12).

# The cross section of high energy muons

Jan Soedingrekso  
Experimentelle Physik 5b  
Technische Universität Dortmund  
jan.soedingrekso@tu-dortmund.de

This work aims at a more detailed understanding of high energy muons. At muon energies between TeV and PeV, the relevant range for Neutrino Astronomy Experiments, the muon cross sections have uncertainties of several percent. Since these uncertainties worsen the reconstruction of these particles, the goal is to reduce the uncertainty to below percent level. The approach is divided into three parts. In a first theoretical step, corrections to the bremsstrahlung and pair production cross section are calculated, since these processes are the most dominant ones. For the second simulation step, the improved cross sections are implemented into the Monte-Carlo Simulation leading to a complete restructuring of the code. The final analyzing step deals with the measurement of cross section normalization. In this report the concept for the final step to fit the bremsstrahlung normalization just using the simulation library PROPOSAL is shown.

## 1 Introduction

High energy muons are the dominant event type, measured by most underground detectors from proton decay experiments to neutrino observatories. To reconstruct these muons, large amounts of Monte-Carlo Simulations representing the measured data are necessary. The resolution of the event reconstruction is therefore limited by the precision of the Monte-Carlo Simulation. To further improve the muon reconstruction, more precise Simulations are needed. In the IceCube Experiment [2] the muons are propagated with the Lepton propagator PROPOSAL [5]. One advantage of PROPOSAL are the multiple cross section parametrizations, that are implemented to do systematic studies on the effects of different muon cross sections. In a new version of PROPOSAL [3, 6] a

complete restructuring of the code was done and a simple exchange of the cross section calculation is now provided. With this new version, a study of the muon cross section was conducted.

## 2 Energy loss Spectrum

A measurement of the muon cross section can be done with the distribution of energy losses. Having a sample of single muons with the same energy and traveling the same length through the same medium, the energy spectrum of the secondaries are representing the cross sections, if the muon energy remains nearly constant. In reality the muons have energies underlying a power-law spectrum, traveling different lengths through the detector, while losing their energy and in case of atmospheric muons, might arrive in bundles.

Instead of atmospheric muons, neutrino-induced muons can be used, which only occur as single muons. A bundle can also be created by single muons via the production of muon pairs. Although the energy loss through those processes can be neglected [4], the additional energy losses through these muons influence the secondary distribution. Unfortunately the effect is not considered in this study, because this process is not yet implemented in PROPOSAL.

The different lengths only increase or decrease the amount of secondaries in the detector, so the secondary spectrum can be normed in length. The power law spectrum of the muon energies is smearing out the secondary distribution. Another problem with the muon energies is, that the muons lose their energy while propagating through the detector, so the later secondaries are produced by muons with lower energies compared to the losses when the muon enters the detector. If the detector consist of a cubic kilometer of ice and the propagated length through the detector varies between 100 m and 1 km, the difference in the muon energies between entering and leaving the detector is in most cases negligible. However, if large energy losses of around a tenth of the muon energy or more occur, it is not valid anymore for the later losses, to correspond them to the initial muon energy. Therefore the analysis is more sensitive to higher energy losses than to small losses and instead of measuring the complete cross section, only the bremsstrahlung normalization is analyzed.

To study the effect of these assumptions and approximations, a study with the Lepton-propagator PROPOSAL was made, propagating  $10^5$  muons with energies between 10 and 30 TeV and a spectral index of  $-3$  through 100 m to 1 km of ice. The energy distribution of the secondaries is shown in figure 1. Here, the secondaries consist of the four interaction processes Ionization, pair production, bremsstrahlung and inelastic nuclear interaction, as well as the electrons originating from the muon decay, which are also visible in the detector.

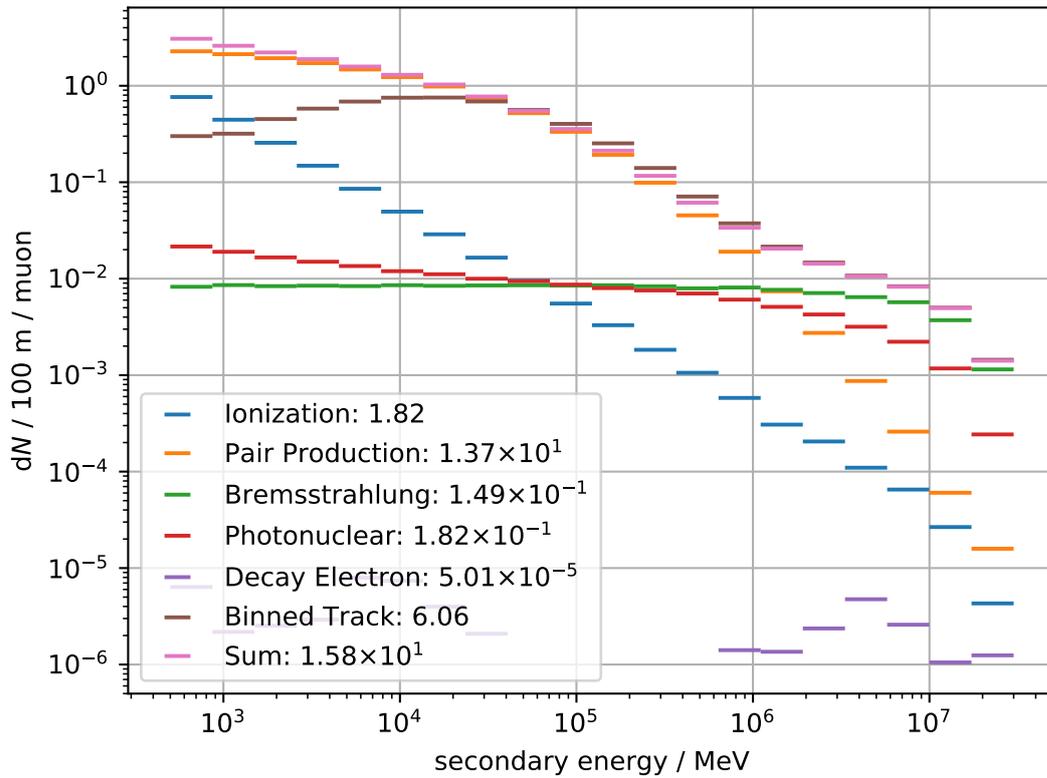


Figure 1: Energies of the different secondaries when propagating  $10^5$  muons with energies between 10 TeV and 30 TeV through ice with lengths between 100 m and 1 km. Also the sum of the energy losses in 15 meter bins is shown. The number after the process in the legend representing the amount of interactions in 100 m for the specific process.

IceCube consists of around 5000 Photodetectors at 86 Strings with a String spacing of about 120 meter. These Photodetectors are deployed into the glacial ice at the geographic south pole in 1500 to 2500 meter depths, constructing an instrumented volume of  $1 \text{ km}^3$ . This sparse DOM-density, good for detecting rare high energy neutrino events, weakens the resolution of single energy losses. The most challenging part for the main analysis will be the poor detector resolution, so that just high energy losses can be well reconstructed. The reconstruction mechanism [1] used in IceCube bins the muon track into 15 meter bins and fits the energy loss for each bin, which is indicated by the Binned Track in figure 1.

### 3 Bremsstrahlung Normalisation Measurement

Due to the different muon energies during the muon propagation and the binned energy loss reconstruction, only the bremsstrahlung normalization is fitted. For that, a secondary spectrum like in figure 1 is created for different bremsstrahlung multipliers between 0.5 and 1.5, shifting the cross section up and down. The difference in each energy loss bin is then parametrized with a linear function and feed into a poisson likelihood fit for each bin. The performance of this study is then tested while injecting secondary spectrums created with random multiplier and fitting the normalization. For 40 randomly sampled multipliers, the fit using the secondary spectrum recovers the injected multiplier with a mean deviation of 1.0(8) %.

### References

- [1] Mark G. Aartsen et al. “Energy reconstruction methods in the IceCube neutrino telescope”. In: *Journal of Instrumentation* 9 (2014), P03009. DOI: 10.1088/1748-0221/9/03/P03009.
- [2] Abraham Achterberg et al., The IceCube Collaboration. “First year performance of the IceCube neutrino telescope”. In: *Astroparticle Physics* 26.3 (2006), pp. 155–173. DOI: 10.1016/j.astropartphys.2006.06.007. arXiv: astro-ph/0604450.
- [3] Mario Dunsch, Jan Soedingrekso, et al. *Recent Improvements for the Lepton Propagator PROPOSAL*. 2018. arXiv: 1809.07740.
- [4] S.R. Kelner, R.P. Kokoulin, and A.A. Petrukhin. “Direct Production of Muon Pairs by High-Energy Muons”. In: *Physics of Atomic Nuclei* 63 (2000), pp. 1603–1611. DOI: 10.1134/1.1312894.
- [5] Jan-Hendrick Koehne et al. “PROPOSAL: A tool for propagation of charged leptons”. In: *Computer Physics Communications* 184 (2013), pp. 2070–2090. DOI: 10.1016/j.cpc.2013.04.001.
- [6] Jan Soedingrekso. *Improvements of the Lepton Propagator PROPOSAL*. Technischer Report für das SFB 876 Gradiertenkolleg. TU Dortmund University, 2017. URL: [https://sfb876.tu-dortmund.de/PublicationFiles/sfbgk\\_etal\\_2017b.pdf](https://sfb876.tu-dortmund.de/PublicationFiles/sfbgk_etal_2017b.pdf).





## Subproject C5

# Real-Time Analysis and Storage of High-Volume Data in Particle Physics

Bernhard Spaan

Jens Teubner

# Efficient Parallel Processing of High-Volume Scientific Data Streams using Event-complexity-driven Machine Placement (EcoMaP)

Thomas Lindemann

Databases and Information Systems Group (DBIS)

TU Dortmund University

thomas.lindemann@cs.tu-dortmund.de

The main objective of our research is to get over today's hardware restrictions in processing capabilities due to electrical power consumption or thermal discharge constraints. We are working in particular on handling large continuous event data streams, which are characterized by a enormous heterogeneity of the event complexity. One of our approaches is to use modern hardware of different architectures to aim the event heterogeneity and get the best results in performance and energy consumption. To validate our approaches, we are processing real world scientific use cases on our self-developed experimental systems, that we get from our collaboration partners of the particle physics department in the SFB876 collaboration, while metering the execution time and energy consumption compared to existing solutions.

Since the last report, we did analysis on the efficiency of Tracking Algorithms used in the LHCb software framework. Our ongoing work is about an Event-complexity-driven Machine Placement Algorithm, briefly referred as *EcoMap*, which is able to place the events on the most suitable hardware with focus on energy, processing time and machine load. Our heterogeneous hardware cluster consists of high-performance Intel Xeon Cores and power-efficient ARM Cortex cores. Using placement strategies based on the event complexity, we try to minimize the computation time per event on the one hand, but the energy on the other hand. Usually, these two characteristics behave in contrary ways and thus, to combine these two measures, we have chosen the energy-delay-product as an efficiency metric.

# 1 Introduction

The LHCb project at CERN is a large and complex research project grown over the last decades. Its general scope is to explain the matter/anti-matter asymmetry. At the time of writing this report, we are working on this project for three years. Our main topic in the SFB876 is the C5 Sub-Project, where a continuous stream of hits inside the LHC (Large Hadron Collider) is produced by the several stages of the LHCb detector, which have to be processed in real time, since there are no capabilities to store all collision events permanently with the current storage technology. Thus, the HLT (High Level Trigger) needs to select events that have to be stored for further analysis. [1] Our specific research topic is to improve the High Level Trigger (HLT) decision time and explore new techniques to process all the experiment data in hard time constraints and sample it down to the maximum load which the storage can handle.

Ressource	Xeon Server	Odroid C2 Cluster
CPU Nr.	2	40
CPU Freq.	2.80 GHz	1.50 GHz
Cores	24	160
Threads	48	160
RAM	256 GB	80 GB
Power	203 W*	120 W <sup>+</sup>

\*CPUs only by perf stat

<sup>+</sup>Secondary power measurement

Table 1: Characteristics of our Test Systems.

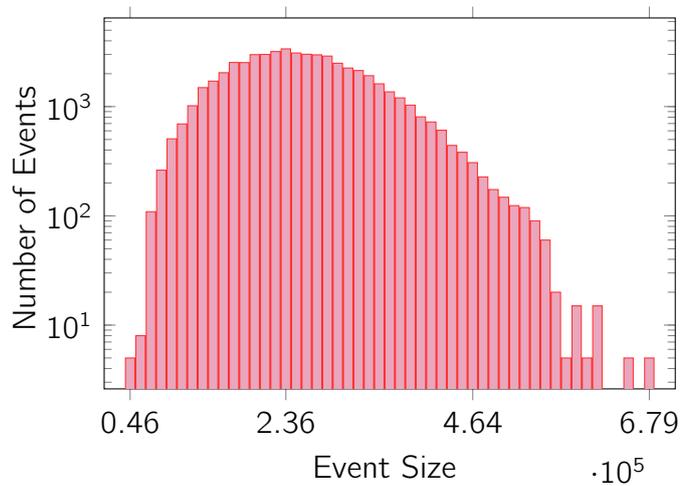


Figure 1: Test Data Set Event Size Histogram

We used some very different test systems as shown in Table 1, on the one hand a dual socket Intel Xeon E5-2695 Server with 48 virtual cores, which represents state of the art hardware, and on the other hand a self constructed ARM cluster with 160 ARM Cortex-A53 cores. The single core performance of the Cortes-A53 cores is of course lower than the Xeon cores, but these architecture provides more efficiency and therefore more cores at the same level of energy, package size and thermal discharge.

The goal is to find new solutions for processing this big amounts of data with limited resources faster than it has been performed in the first run of the LHCb project and allow the physicists to make experiments with more precise decisions. A challenge of the data is the wide spread of heterogeneity of the events, which is caused by different interactions of the particles after the collision. The event size histogram in Figure 1 shows the event sizes on a logarithmic scale.

We have the approach not to appraise the event heterogeneity as a problem, but rather as a part of the solution, because it makes it possible to distribute events based on their complexity on the most suitable hardware.

## 2 Event complexity-driven Machine Placement

In the current configuration, the High Level Triggers concept used by the LHCb Project is a large computing grid of state of the art Xeon-based server machines.

Due to the heterogeneous event complexity as investigated in Figure 1, our ongoing work is about an Event-complexity-driven Machine Placement Algorithm, briefly referred as *EcoMap*, which is able to place the events on the most suitable hardware with focus on energy, processing time and machine load.

In Figure 2, one can see that the overall energy consumption is at the lowest level when processing events on ARM Cortex-A53 Cores (highest threshold), while on the other hand the execution time effort is lower when placing all events on state of the art Intel Xeon E5-2695 cores (lowest threshold). One might think that the race is already done by the ARM cluster, because of the lower energy amount needed and the fact that the energy is product of power and time consumption. But unfortunately, it's not that easy, because there are still strict time constraints for the average processing time of a single event, due to the constant event stream which has to be processed in time.

As a better metric, we propose to use the energy-delay-product (EDP) as a metric for system efficiency, because it aims the problem to get a 'fair' trade off between energy on the one hand, but on the other hand it further has the advantage of perfect scalability between energy consumption and execution time. For example, if a system's EDP is fine, but the processing time is too high, the number on systems can be doubled, which leads to a half of the processing time and two times of the energy costs, so that the EPD is constant. The approach is to place the events on different hardware architectures, which are the most suitable for the job at the moment of processing to minimize the EDP. Thus, we now took a closer look to the specific relation between event complexity and processing characteristics, instead of just looking at overall energy or time for the whole workload and tried to find criteria for making an EDP optimizing placement decision.

In our placement experiments , we started with processing all events on the Xeon Server and then we slowly raised the threshold, which means we moved all events below the current threshold to the ARM-Cluster.

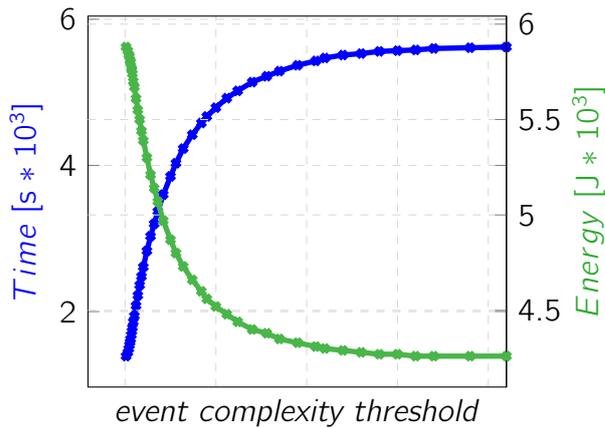


Figure 2: Time & Energy Consumption

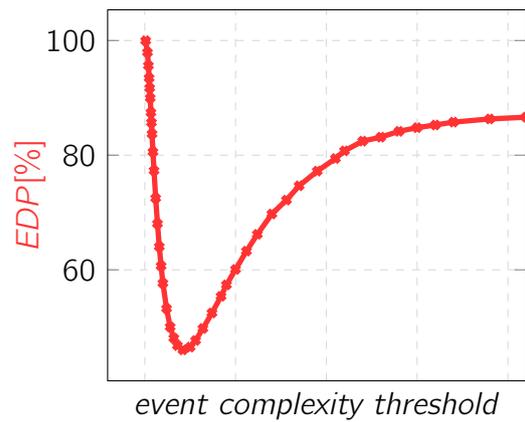


Figure 3: Energy-Delay-Product

An interesting observation is that there is a minimum in the EDP, which is not at the outer right end of the plot, where all events are placed on the ARM-Cluster. Instead, regarding to the EDP as control factor for placement, there is an optimal threshold and it is reasonable to place small events on the Xeon Server and bigger events on the ARM Cluster to reach this optimum.

We can also see that the EDP in the minimum is about 45% of the EDP compared to placing all events on the server, which is marked as 100%. This is significantly lower and means that it is possible to save an very big amount of energy and time.

### 3 Conclusion and Future Work

Our experiments have shown that the ARM Cortex-A53 low power cluster has got a better energy efficiency than a state-of-the-art Xeon dual socket server system, while the Xeon Server has a better single thread performance. Moreover, we could show when adding a dynamic placement algorithm like EcoMaP to use cases with event streams of heterogeneous event complexities, there is still optimization potential in the balance of energy and time consumption by placing the events at runtime to the most suitable hardware.

In addition, our research is further aiming a distributed storage solution after the Trigger, which offers fast OLAP functionality on stored data, but also OLTP capabilities during data acquisition.

### References

- [1] The LHCb collaboration. Technical Report LHCb Tracker Upgrade Technical Design Report. *CERN-LHCC-2014-001. LHCb-TDR-015*, Feb 2014.

# Improved precision measurement of *CP* violation

Margarete Schellenberg  
Lehrstuhl für Experimentelle Physik 5  
Technische Universität Dortmund  
margarete.schellenberg@tu-dortmund.de

One of the four big experiments at the Large Hadron Collider (LHC) near Geneva is the LHCb experiment [1]. Its main focus is the research of the asymmetry between matter and anti-matter in the observable universe. Matter and anti matter should have been produced in equal amounts during the early stages of the universe. Today we observe a large asymmetry between the two, so that it is assumable, that physical laws influence matter and antimatter in different ways. Physicists at the LHCb experiment are investigating the charge-parity (*CP*) violation in decays of beauty and charm hadrons as one necessity cause for this asymmetry. Due to the focus on hadrons containing *b* and *c* quarks, the LHCb detector is designed as a single-arm forward spectrometer (see Figure 1). At the interaction point of the proton beams, which lies inside the vertex locator (VELO), a large amount of particles is produced by many different physical processes. These particles decay into new particles, which fly through the detector and interact with the detector material. Different subdetectors are responsible for the reconstruction and identification of these particles. The flight path reconstruction is performed by the tracking system consisting of the VELO, TT, IT and OT. The particle identification system is composed of the two RICH detectors, the two calorimeters ECAL and HCAL, as well as the Muon chambers. Information from all components is used to reconstruct the complete decay chain, by combining the tracks and tracing them back to heavier mother particles.

At the LHC, collisions are produced at a rate of 40 million collisions per second. Considering down times of the collider, the experiment has to handle nearly  $40 \times 10^{14}$  collisions per year. It is not possible to store every collision, as each one of them creates about 100 kB of data. Thus, an online trigger system is utilised, which is passed by only a few

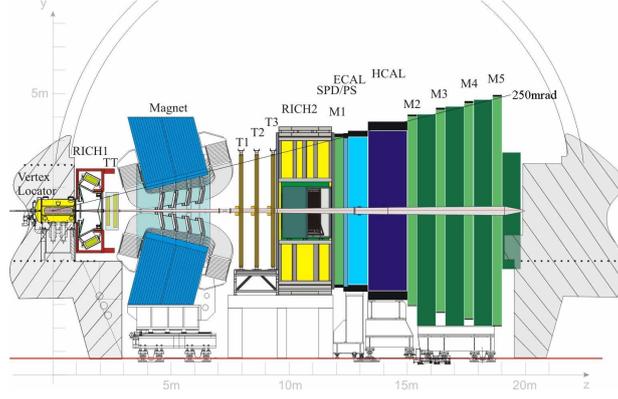


Figure 1: Scheme of the LHCb detector, illustrating the various subdetectors for reconstruction and identification of particles and their tracks. [2]

per mille of the events, resulting in an amount of data that is savable on a large storage cluster. After a centralised loose preselection of the complete dataset, the data is used by physicists, who store the data that corresponds to their analysis' conditions in form of ROOT [4] nTuple structures. These nTuples are much smaller than the full recorded dataset, but can still reach a size of several hundreds of Gigabytes.

To show the necessity of handling the data in an efficient way, a measurement of time-dependent  $CP$  violation in the decay  $B^0 \rightarrow D^{*\pm} D^\mp$  is being performed. The reconstruction is done with  $D^- \rightarrow K^+ \pi^- \pi^-$  and  $D^{*+} \rightarrow D^0 \pi^+$ , where the  $D^0$  decays into  $K^- \pi^+$ .

In  $B^0 \rightarrow D^{*\pm} D^\mp$  decays a a decay-time-dependent  $CP$  asymmetry can be measured, because freely propagating  $B^0$  mesons can mix into their anti particle state ( $\bar{B}^0$ ) and vice versa and the charge-conjugated final states  $D^{*+} D^-$  and  $D^{*-} D^+$  are common to  $B^0$  and  $\bar{B}^0$  mesons. The asymmetry results from the interference between the amplitudes of the direct decays and decays after  $B^0$ - $\bar{B}^0$  mixing:

$$A_f(t) = \frac{\Gamma(\bar{B}^0(t) \rightarrow f) - \Gamma(B^0(t) \rightarrow f)}{\Gamma(\bar{B}^0(t) \rightarrow f) + \Gamma(B^0(t) \rightarrow f)} = \frac{S_f \sin(\Delta m t) - C_f \cos(\Delta m t)}{\cosh\left(\frac{\Delta \Gamma t}{2}\right) + D_f \sinh\left(\frac{\Delta \Gamma t}{2}\right)}. \quad (1)$$

The decay-time-dependent asymmetry is given by the difference between the time-dependent-decay widths of  $B^0$  and  $\bar{B}^0$  mesons decaying into the final state  $f$ , normalised to the sum. An analogous asymmetry exists for the final state  $\bar{f}$ .  $B^0(t)$  and  $\bar{B}^0(t)$  denote the initial  $B$  flavour and the parameters  $\Delta m$  and  $\Delta \Gamma$  are the differences of the masses and decay widths between the heavy and light mass eigenstates concerning the  $B^0$ - $\bar{B}^0$  system [6]. With some assumptions, the time-dependent asymmetries become

$$A_f(t) = S_f \sin(\Delta m t) - C_f \cos(\Delta m t), \quad A_{\bar{f}}(t) = S_{\bar{f}} \sin(\Delta m t) - C_{\bar{f}} \cos(\Delta m t), \quad (2)$$

where  $S_f$ ,  $S_{\bar{f}}$ ,  $C_f$  and  $C_{\bar{f}}$  are the  $CP$  observables.

The analysis of  $CP$  violation in  $B^0 \rightarrow D^{*\pm} D^\mp$  decays can be divided into two parts. The first comprises the selection, where signal and background decays are separated. Due to wrong reconstruction or misidentification, these so-called background decays can be mistakenly be present in the data set and hence, need to be removed in the selection. The selection chain starts with rectangular cuts on kinematical requirements. In this context a rectangular cut requires a certain measured quantities to have values in a specific range. Theses ranges are well-defined and do not vary depending on other quantities in the event. At a latter stage in the selection a multivariate analysis is performed, which exploits correlations between quantities. Although good control over the different backgrounds is achieved, background contributions are still present in the data set at the end of the selection. Thus, a statistical background subtraction is performed by fitting the invariant  $B^0$  mass.

The second part of the analysis is a maximum-likelihood fit of the  $B^0$  meson decay-time distribution to measure the  $CP$  parameters. Figure 2 shows a representative fit of the decay-time distribution using simulated data. As can be seen in Fig. 2, the decay-time distribution does not follow a usual exponential distribution, but shows a drop towards lower values. This effect originates from inefficiencies in the reconstruction that are more likely at low decay times. To describe the deviation of the decay-time distribution from the exponential shape, a parametrisation including cubic splines is used [5]. Another

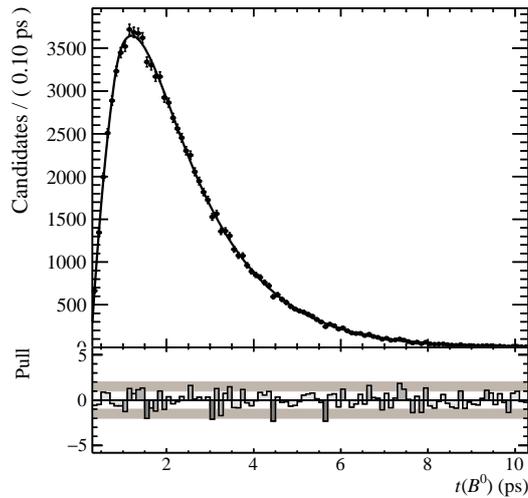


Figure 2: Decay time distribution of the decay  $B^0 \rightarrow D^{*\pm} D^\mp$  of simulated data.

effect that needs to be taken into account is the decay-time resolution, which results from the limited precision in the determination of momenta and vertex positions. In order to describe the distribution of the reconstructed decay time appropriately, this effect can be caught up by convolving the theoretical decay-time PDF with a resolution model describing the probability density of the difference between the true and reconstructed decay times. As presented in Fig. 3 this deviation can be described by the sum of three

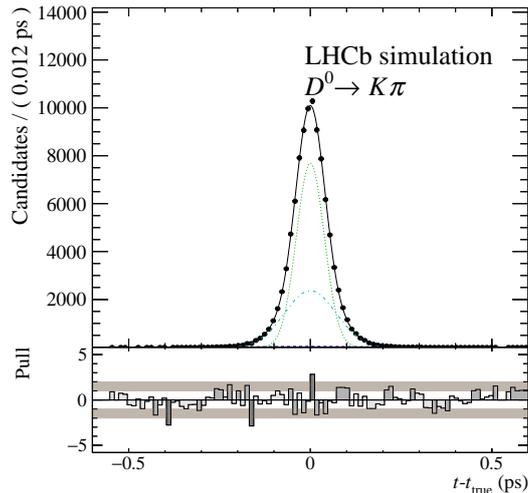


Figure 3: Difference between true and reconstructed decay times of simulated data.

Gaussians. Results of the data fit cannot be shown yet, because the analysis is ongoing and the central values of the  $CP$  parameters are still unknown. However, preliminary statistical uncertainties were already obtained. For the parameter  $S$  this uncertainty is about twice as large as in the respective analyses of the Belle and BaBar experiments [3,7]. Thus, to become competitive with the  $B$ -factories, the statistics needs to be increased.

## References

- [1] A. A. Alves, Jr. et al. The LHCb detector at the LHC. *JINST*, 3:S08005, 2008.
- [2] R Aaij et al. Letter of Intent for the LHCb Upgrade. Technical Report CERN-LHCC-2011-001. LHCC-I-018, CERN, Geneva, Mar 2011.
- [3] Bernard Aubert et al. Measurements of time-dependent  $CP$  asymmetries in  $B^0 \rightarrow D^{(*)+}D^{(*)-}$  decays. *Phys. Rev.*, D79:032002, 2009.
- [4] R. Brun and F. Rademakers. ROOT: An object oriented data analysis framework. *Nucl. Instrum. Meth.*, A389:81, 1997.
- [5] Carl de Boor. *A practical guide to splines*, volume 27 of *Applied Mathematical Sciences*. Springer-Verlag New York, 1978.
- [6] K. A. Olive et al. Review of particle physics. *Chin. Phys.*, C38:090001, 2014.
- [7] M. Röhrken et al. Measurements of Branching Fractions and Time-dependent  $CP$  Violating Asymmetries in  $B^0 \rightarrow D^{(*)\pm}D^\mp$  Decays. *Phys. Rev.*, D85:091106, 2012.

# GPU accelerators in the future

## LHCb-Triggersystem

Holger Stevens

Lehrstuhl für Experimentelle Physik 5

Technische Universität Dortmund

holger.stevens@tu-dortmund.de

The LHCb experiment entered an upgrade phase in Dezember 2018, which will last 2 years. Not only the detector but also the computing farm will be changed. The planning for this started 10 years ago. A certain amount of computing power was planned to be bought by a fixed budget, but meanwhile "Moore's Law died" so the plan cannot be converted. There are different solution approaches in the LHCb collaboration to reach the necessary computing power, one is the usage of GPUs.

Computing tasks can only benefit from the multi processor architecture of a GPU, if they are massively parallelizable. Different algorithms are running in the online farm to reconstruct the events. Most of them are independent and parallelizable, On the one hand it is the combination of hits to reconstruct the tracks and on the other hand the processing of RAW informations.

Last year the report was focused on the development of new tracking algorithms. This report will describe the decoding of the so called RAW banks of the SciFi tracker. These contain all informations of the channels which detect a signal. The known channel IDs need to be translated to physical positions in the detector.

# 1 Introduction

The LHCb experiment is one of the four big experiments located at the Large Hadron Collider (LHC) near Geneva, Switzerland. Its main focus is the search for rare decays and effects of  $CP$ -violation in decays of beauty and charm hadrons [1]. Due to some physical constraints in the production of  $b$  and  $c$  quarks through proton proton collisions the LHCb detector is designed as a single-arm forward spectrometer. Over the past years the understanding of the detector and its systematical effects has reached an almost perfect level. At the moment, the most limiting factor for analyses is the statistical uncertainty. The only way to improve this is to massively increase the dataset. For this reason, as mentioned before an upgrade of the experiment is foreseen [5]. Most of the

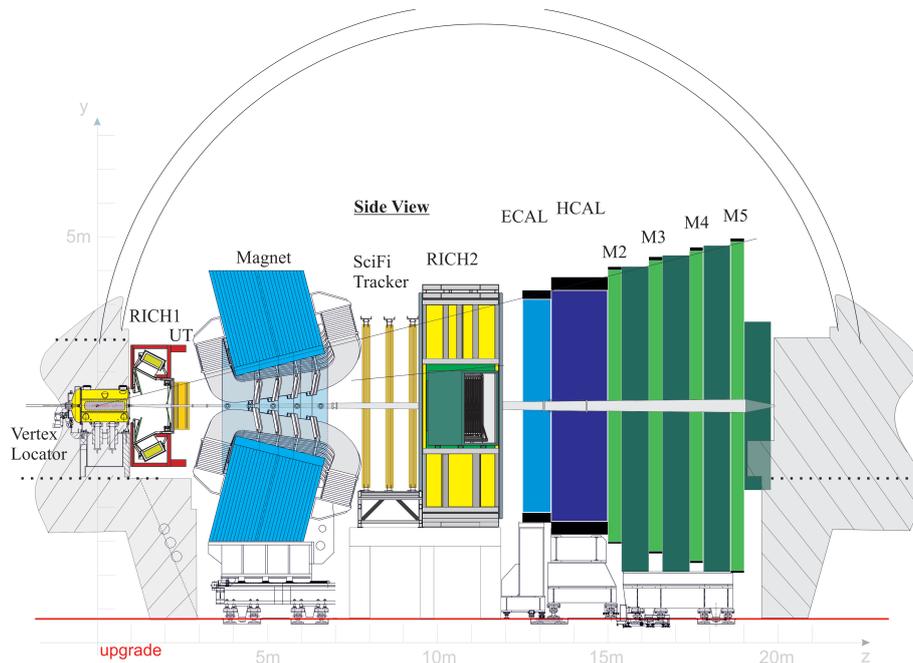


Figure 1: The LHCb upgrade detector with the various subdetectors for the identification of particles and reconstruction of their tracks [2].

existing detector will be replaced, but the general structure of the detector will remain the same. In the Vertex Locator (Velo) the position of the primary interaction is detected. The Upstream Tracker (UT) and the SciFi Tracker also belong to the tracking system. Other components like the Ring Imaging Cherenkov Detectors (RICH), the Electronic Calorimeter (ECAL), the Hadronic Calorimeter (HCAL) and the Muon Chambers (M2-M5) are used for the particle identification.

A major change is the new triggerless readout of the full detector with 40 MHz. This leads to a data rate of 40 Tb/s which is the input of the high level trigger (HLT) [6].

## 2 Details of the SciFi Tracker

Before the actual decoding part is described, some technical details of the SciFi Tracker are explained to understand the data format.

In total the Tracker consists of 3 stations with 4 layers each. These layers are build out of so called modules, which again are composed of 8 fibre mats. Each mat is connected to 4 SiPM<sup>1</sup> arrays. These are connected to the so called FrontEnd electronics (FE). A FPGA based algorithm is looking for clusters of hits. A real signal has a different width a high than a background hit, caused for example by thermal noise. Only the information from found clusters, so called zero suppressed data, is send from the FE to the BE (BackEnd electronics). The BE has to pack the data from several FE boards and send it to the computing farm where the actual decoding will take place. It is important to mention that the stations have a different widths, T1 and T2 consist of 10 modules and T3 is larger with its 12 modules.

## 3 Status and outlook

The aim of one working group of the LHCb collaboration is to run the first Trigger stage (HLT1) on GPU [4]. The project, Allen, can be found at <https://gitlab.cern.ch/lhcb-parallelization/Allen>. Multiple streams can send data to the GPU. Than a visitor service will load the necessary Algorithms. One of these is the SciFi hit decoding. Every data-package has a global header which contain information about the associated proton-proton collision, like a time stamp. This information is followed by detector component specific informations, for the SciFi this are the station number, MatID etc. with these information a position in the detector is unambiguously defined. The actual numbering scheme is constant for a certain time and can be a loud as a constant in the code [7]. During the development of the decoder the data-format was changed to reduce the needed bandwidth. The amount of Hits is not longer in the header [3]. To allocate the correct size of memory for the result vector of decoded Hits a pre count is necessary.

Successfully a GPU version of the SciFi decoder was implemented. There will be further improvements and fine tuning of the code, but first the input data from the encoder has to be fixed. As mentioned before the global header was changed, but also further changes are discussed at the moment. For example the numbering scheme of the detector. This has a direct effect of the procedure who the Hits are mapped to Threads and Warps. At the moment the biggest problem is the fixed number of RAW banks per station. In the first two stations a RAW bank is connected to a single Module, but in the last station

---

<sup>1</sup>silicon photonmultiplier

this is not possible. This leads to a ugly branching in the code, which slows down the speed and efficiency of the algorithm significant.

## References

- [1] Letter of Intent for the LHCb Upgrade. Technical Report CERN-LHCC-2011-001. LHCC-I-018, CERN, Geneva, Mar 2011.
- [2] LHCb Collaboration. LHCb Tracker Upgrade Technical Design Report. Technical Report CERN-LHCC-2014-001. LHCB-TDR-015, Feb 2014.
- [3] Sevda Esen, Jeroen Van Tilburg, Luigi Del Buono, Pierre Billoir, and Louis Henry. Clustering and rawbank decoding for the SciFi detector. Technical Report LHCb-INT-2018-024. CERN-LHCb-INT-2018-024, CERN, Geneva, Jul 2018.
- [4] Stefano Gallorini, Donatella Lucchesi, Alessio Gianelle, Silvia Amerio, and Marco Corvo. First experiences with a parallel architecture testbed in the LHCb trigger system. Technical Report LHCb-PUB-2017-015. CERN-LHCb-PUB-2017-015. 3, CERN, Geneva, Apr 2017.
- [5] Christian Joram. LHCb Scintillating Fibre Tracker Engineering Design Review Report: Fibres, Mats and Modules. Technical Report LHCb-PUB-2015-008. CERN-LHCb-PUB-2015-008, CERN, Geneva, Mar 2015.
- [6] CERN (Meyrin) LHCb Collaboration. Computing Model of the Upgrade LHCb experiment. Technical Report CERN-LHCC-2018-014. LHCB-TDR-018, CERN, Geneva, May 2018.
- [7] Jeroen Van Tilburg. SciFi readout numbering scheme. Technical Report LHCb-INT-2016-044. CERN-LHCb-INT-2016-044, CERN, Geneva, Nov 2016.