Institute of Software Engineering and Artificial Intelligence

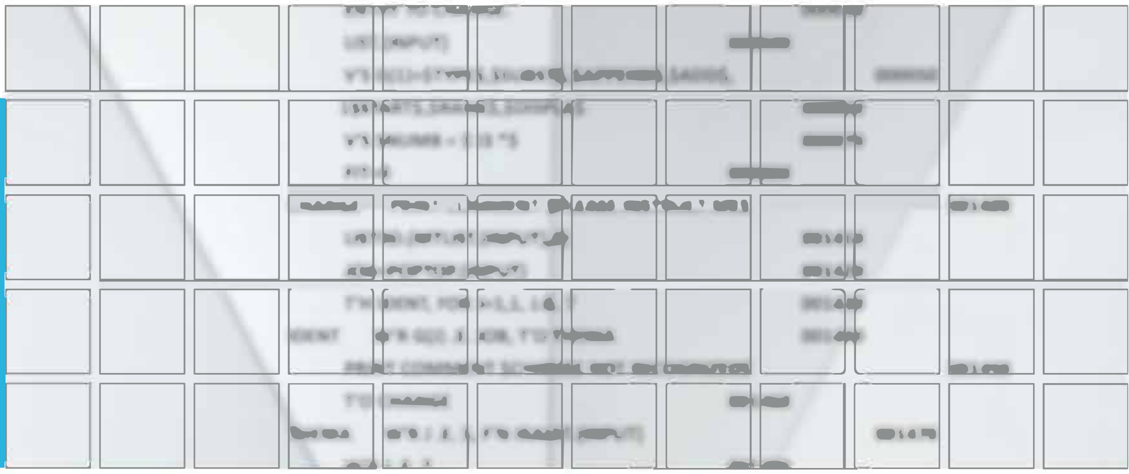

The TU Graz Institute of Software Technology

The First 20 Years

The TU Graz Institute of Software Technology

The First 20 Years

Franz Wotawa (Editor)
TU Graz, Institute of Software Engineering and Artificial Intelligence (SAI)
Graz, Austria

Verlag der Technischen Universität Graz

Editor: Franz Wotawa
Cover picture: Franz Wotawa
Print: Buchschmiede (DATAFORM Media GmbH)

2026 Verlag der Technischen Universität Graz
Technikerstraße 4, 8010 Graz
verlag@tugraz.at
www.tugraz-verlag.at

ISBN (print) 978-3-85125-977-3
ISBN (e-book) 978-3-85125-978-0
DOI 10.3217/978-3-85125-977-3

Preface

The Institute of Software Technology (IST) was founded 20 years ago in the year 2003. 20 years of research and teaching in software engineering, artificial intelligence, and theoretical computer science is the best reason to summarize what has been done. Likewise, 20 years is also a good time frame to reminisce about what worked well and where there is potential for improvement. In this book, the members of the IST present their research activities, history, statistics, and benchmarks indicating the institute's performance after 20 years. To make it short, the IST performed excellently, exceeding expectations. I hope that readers enjoy this book. The content is very diverse, like the institute itself, spanning a huge research area and conducting research from curiosity-driven to applied. As the head of the IST and the editor of this book, I want to thank all contributors, namely (in alphabetical order) Oswin Aichholzer, Bernhard Aichernig, Alexander Felfernig, Yannic Maus, Gerald Schweiger, Wolfgang Slany, Gerald Steinbauer-Wagner, Birgit Vogtenhuber, and Johannes P. Wallner for their texts and other inputs.

Note that with 2025 the faculty and also the IST implemented several changes. First, all colleagues, i.e., Oswin Aichholzer, Birgit Vogtenhuber, and Yannic Maus, working in the area of theoretical computer science joined the newly founded Institute of Algorithms and Theory. Second, the name of the IST changed to Software Engineering and Artificial Intelligence (SAI) to reflect the research content of the remaining members of the institute.

Franz Wotawa (Head of the Institute of Software Engineering and Artificial Intelligence), Graz, December 2025

Table of Contents

Chapter 1

Introduction

The Institute of Software Technology (IST) at the Graz University of Technology was founded in 2003 as an independent unit. Before, it was part of Prof. Hermann Maurer's IICM as a working group. Since its beginnings, the IST has always carried out research in Software Engineering, particularly Formal Models focusing on Quality Assurance Methodology. Research on methodologies for assuring and improving software quality is still IST's largest and most important part. However, the IST extended its research activities to cover Theoretical Computer Science, Robotics, Artificial Intelligence, and Software Engineering, always aiming at research excellence and impacting the corresponding research communities. More information regarding past and current research activities can be found in the research chapter of this book.

Besides research, the IST has always been actively involved in teaching. Institute members give lectures in all computer science curricula at all levels, from Bachelor to Master and PhD. In addition, we have contributed substantially to the development of the curricula by serving as Deans of Studies or as members of curricula committees. To illustrate the teaching effort provided by the institute, let us have a look at some figures. The IST ranges in second place regarding teaching effort in total and the number of teaching hours provided each semester. When considering the available potential, i.e., the number of people working at the IST, we are in the first place compared to the other computer science institutes. To keep the effort within good boundaries, further personnel support must be provided, or lectures will not be carried out anymore. Anyhow, despite the big engagement in teaching, the IST also contributes substantially to the research of TU Graz and the Faculty for Computer Science and Biomedical Engineering in particular. Further quantitative and qualitative data will be provided in Chapter 5.

The IST currently comprises three full professors (Alexander Felfernig, Wolfgang Slany, and Franz Wotawa), four associate professors (Bernhard Aichernig, Oswin Aichholzer, Gerald Steinbauer-Wagner, and Birgit Vogtenhuber), two assistant professors on a tenure track (Yannic Maus and Johannes P. Wallner), two senior lecturers (Roxane Koitz and Martin Stettinger), and ten Post-Docs and PhDs funded by TU Graz for supporting research and teaching activities. In addition, the IST has two senior post-docs (Birgit Hofer and Gerald Schweiger), one senior scientist (Iulia Nica), another 32 researchers (mainly PhDs), and 9 student project members externally funded and supporting carrying out the IST research projects.

The backbone of the institute supporting all organizational activities are our secretaries (Elisabeth Orthofer and Petra Schindler) and our computer administration (Jörg Baumann and Mario Wellik). Without their support, effective work is hardly possible. Hence, the IST employs 60 persons for administration, research, and teaching. This is a tremendous increase compared with the 10 persons working at the IST in its founding year. The institute is informally structured in research groups. Each research group has a senior post-doc or a professor as the head. For more details regarding the different groups' research activities, look at this book's research chapter.

The book is organized as follows: First, we present the program and the talks that were given at the colloquium 20 Years IST, which was carried out on December 15, 2023, in lecture hall i12. The institute invited friends and guests to contribute to the celebrations. There was a welcome reception at the Aula one day earlier and a party afterward at the Mensa Inffeldgasse. Afterward, we briefly summarize the history of the institute from its origins. The research chapter comprises contributions from all research groups of IST. The head of the respective groups present their objectives and their current activities. If someone is interested in having a picture of the research interests, problems, and solutions, reading this chapter is a must. In particular, the following group presentations are available that show the large diversity of research interests at the institute:

- Formal Development Methods (FDM) – Bernhard Aichernig
- Discrete and Computational Geometry (DCG) – Oswin Aichholzer
- Applied Software Engineering & AI (ASE) – Alexander Felfernig
- Algorithms and Complexity (AC) – Yannic Maus
- Intelligent Energy Systems (IES) – Gerald Schweiger
- Agile Software Development, Mobile Applications, and Didactics of Programming (ASD) – Wolfgang Slany
- Autonomous Intelligent Systems (AIS) – Gerald Steinbauer-Wagner
- Discrete Mathematics & Theoretical Computer Science (DMTCS) – Birgit Vogtenhuber
- Knowledge Representation (KR) – Johannes P. Wallner
- Software Engineering and Artificial Intelligence (SEAI) – Franz Wotawa

Furthermore, we outline important statistics using metrics indicating the effort and output of teaching and research. We further benchmark the outcome with expectations and other institutes of the faculty. Finally, we conclude the book and outline future activities.

Chapter 2

Colloquium 20 Years IST

To celebrate 20 years of IST, we organized a colloquium on the 15th of December 2023 and invited colleagues from Europe who supported us to grow. You find a detailed schedule, short CVs of invited speakers, and the title and abstract of their talks in this chapter.

2.1 Schedule

Place: Lecture hall i12 at the Inffeldgasse Campus (Inffeldgasse 16b, A-8010 Graz).

10:00-10:40	**Welcome**
10:40-12:10	**Session 1**
	Gerhard Friedrich, *The (artificial) engineer – data fitting or knowledge-based reasoning?*
	Emo Welzl, *Random (Order Types of) Point Sets*
	Kim G. Larsen, *Shielding and Verification of Learning-Controlled Cyber-Physical Systems*
12:10-13:10	***Lunch break***
13:10-14:40	**Session 2**
	Forza Cipriano, *Product Configurators: investigations on their application*
	Gordon Fraser, *Has the Search for Tests Come to an End?*
	Louise Travé-Massuyès, *Leveraging the properties of the Christoffel function for anomaly detection in data streams*
14:40-15:00	***Coffee break***
15:00-16:00	**Session 3**
	Juha Tiihonen, *Solutions for configurable products - from research to VariSuite*
	Mihai Nica, *Navigating the Road to AV Reliability: A Comprehensive Look at Testing Strategies in ADAS and Automated Driving Systems*
16:00-16:50	**Alumni talks**
	Michael Reip (incubedIT), Martin Weiglhofer (JungHeinrich), Christop Zehentner (incubedIT), Willibald Krenn (SAL), Stefan Galler (Frequentis)
16:50-17:10	**Closing**

2.2 Invited talks

Prof. Dr. Gerhard Friedrich
University of Klagenfurt

Gerhard Friedrich is a full professor at the University of Klagenfurt, Austria, and heads the Intelligent Systems and Information Systems Research Group. He holds a PhD in computer science from the Vienna University of Technology. His research interests include heuristics in automated problem solving, recommender systems, knowledge representation and processing, logic programming, constraint programming, knowledge acquisition and maintenance, configuration of products and services, automated planning and diagnosis. He has published numerous papers in these areas in prestigious journals. A book on recommender systems published by Cambridge University Press has been translated into Japanese and Chinese. Prof. Friedrich is a Fellow of the European Association for Artificial Intelligence and the Asia-Pacific Artificial Intelligence Association.

Title: The (artificial) engineer – data fitting or knowledge-based reasoning?

Abstract: Currently, the field of artificial intelligence is divided into the machine learning community - dominated by artificial neural networks (ANN) - and the knowledge representation and reasoning (KRR) community. First-principles knowledge is fundamental to solving engineering problems such as the design, configuration, and diagnosis of technical systems. We will demonstrate the power of using first-principles knowledge by means of KRR methods and point out open problems where the integration of ANN and KRR methods will lead to a new generation of artificial intelligence systems.

Prof. Dr. Emo Welzl
ETH Zürich

Emo Welzl has been a Full Professor of Computer Science at the Institute for Theoretical Computer Science of the ETH Zurich since April 1996. His research interests are in foundations of computer science, mainly combinatorial algorithms, in particular computational geometry and applications, combinatorial models for optimization, analysis of geometric structures, randomized methods, and discrete geometry. Career: 1977-1981 Diplom in Applied Mathematics at the Graz University of Technology, Austria. 1983 Ph.D. (supervisor Hermann Maurer) with a topic in formal languages, 1988 habilitation in Foundations of Computer Science, also in Graz. 1984 one-year post-doc at Rijks University Leiden, Netherlands, 1985 visiting professor at the University of Denver, Colorado, USA, for one semester. 1987-1996 Full Professor of Mathematics (theory of computation) at Free University of Berlin. 1994 four-month research stay at the International Computer Science Institute, Berkeley, California, USA. 1991-1996 speaker of the graduate program ”Computational Discrete Mathematics” at Free University, Humboldt University, University of Technology, and the Konrad Zuse Center in Berlin. 2000-2005 speaker (at site Zurich) of the European Graduate Program in ”Combinatorics, Geometry, and Computation”. Grants, prices, and memberships: 1992 Max Planck-Prize with Micha Sharir (Tel Aviv University); 1995 Gottfried Wilhelm Leibniz-Prize; 1998 Fellow, Association for Computing Machinery; 2005 Member of the German Academy of Sciences Leopoldina; 2006 Member of the Academia Europaea; 2007 Member of the Berlin-Brandenburg Academy of Sciences.

Title: Random (Order Types of) Point Sets

Abstract: Roughly speaking, a planar order type is a class of point sets in the plane that cannot be discriminated by a typical convex hull algorithm based on sidedness queries: ”Does point p lie to the left of the oriented line through points q and r”. In this sense, order types are the counterparts of permutations in the context of sorting algorithms. Order types have been intensively investigated in discrete and computational geometry, prominently by the research group in Graz, where - among others - a database of all order types of small point sets has been created and maintained. It is natural to ask for random order types (e.g., for testing convex hull algorithms). We investigate typical properties of uniformly random order types (e.g., what is the expected number of vertices of the convex hull). Moreover, it can be shown that when sampling points uniformly from a disk or a square in the plane, then the resulting order types are concentrated among all order types. In other words, one encounters only a vanishingly small part of all possible inputs to convex hull algorithms. The investigation of properties of random point sets has a long history, dating back to at least 1865 when J. J. Silvester posed his famous Four Point Problem.

Prof. Dr. Kim G Larsen
Aalborg University

Kim is a professor in the Department of Computer Science at Aalborg University within the Distributed, Embedded Systems and Intelligent (DEIS) Unit and director of CISS, Center for Embedded Software Systems. Kim's research interests include semantics, verification and logic, concurrency theory, real-time, embedded, and cyber-physical systems, model checking, and machine learning with numerous applications in transport, energy, and water management. Kim is a prime investigator of the tool UPPAAL, winning the CAV Award in 2013. Kim also holds honorary doctorates from Uppsala U, Cachan, and NEU China. In 2015, Kim won an ERC Advanced Grant with the project LASSO for learning, analysis, synthesis, and optimization of cyber-physical systems. Kim is currently a VILLUM Investigator with the project S4OS: Scalable Analysis and Synthesis of Safe, Secure, and Optimal Strategies for Cyber-Physical Systems.

Title: Shielding and Verification of Learning-Controlled Cyber-Physical Systems

Abstract: Deep reinforcement learning (RL) is increasingly being used in cyber-physical systems (CPS) to achieve unprecedented performance. Successful applications of RL within the control of CPS may be found in traffic systems, medical devices as well as infrastructures for water management and heating of buildings. Despite impressive performance, there are two key drawbacks: (i) the resulting control strategy is difficult to interpret, and (ii) there are no guarantees as to the overall correctness of the system. In this talk, we will promote the use of decision trees as a well-performing, compact, and interpretable surrogate for deep neural network control strategies. Using decision trees, we have extended the well-known notion of shielding to CPS and demonstrated that results in guaranteed correct yet highly optimal control strategies. Furthermore, we have developed the first algorithm to directly verify decision-tree controlled systems in continuous time. The key aspect of our method is exploiting the decision-tree structure to propagate a set-based approximation through the decision nodes.

Prof. Dr. Forza Cipriano
University of Padova

Cipriano Forza holds an M.Sc. degree in Electronic Engineering (with a specialization in Informatics) and a Ph.D. degree in Industrial Innovation Science (Industrial Management) from the University of Padova (Italy). He has been a visiting scholar at Minnesota University, London Business School, and Arizona State University. He is a full professor of Management Engineering. He served as coordinator of the PhD course in Management Engineering and Real Estate Economics at Padova University (from 2007 to 2018). He has been serving as a reviewer for a number of academic journals (JOM, DSJ, IJOPM, IJPE, IJPR, IJTM, Cind, etc.), as a member of the scientific committee of several conferences (EurOMA, MCP-CE, Configuration Workshop, etc.), as associate editor of Journal of Operations Management and as guest editor of JOM and POM. He has widely researched product variety management, product configuration, mass customization, quality management, lean production, product modularization, and IT support for operations. He produced more than 110 publications, 43 of which were in peer-reviewed journals. He served for several years in Research Groups of the Center for Productivity in Veneto and in the OIIV of the Vicenza Chamber of Commerce. He has been on the board of the AiIG (Italian Association of Management Engineering) and EurOMA (European Association of Operations Management). Since 2012, has been serving as a member of the EAJG Ltd Scientific Committee (The International Guide to Academic Journal Quality - formerly ABS list) - The Association of Business Schools (ABS). In 2021, he was awarded the title of EurOMA Fellow for merit in research and service to the discipline of Operations Management.

Title: Product Configurators: investigations on their application

Abstract: Configurators as tools to integrate marketing, new product development, sales, operations, and managerial accounting in contexts where customization and product variety are offered. Insights on how configurators influence quality, time, and ROI performance. The knowledge management side of product configurators. Organizational change and introduction of product configurators. Open research issues in product configuration research: the challenge of developing configurators suitable for SMEs.

Prof. Dr. Gordon Fraser
University of Passau

Gordon Fraser is a full professor in Computer Science at the University of Passau, Germany. He received a PhD in computer science from Graz University of Technology, Austria, in 2007, worked as a post-doc at Saarland University, and was a Senior Lecturer at the University of Sheffield, UK. The central theme of his research is improving software quality, and his recent research concerns the prevention, detection, and removal of defects in software.

Title: Has the Search for Tests Come to an End?

Abstract: Search-based software testing is one of the most popular approaches to automatically generate tests as it is applicable to any testing scenario, easy to implement, but nevertheless effective. Evolutionary search has been successfully applied to test anything from object-oriented code, graphical user interfaces, and autonomous driving systems to games. Despite many success stories, there are remaining fundamental as well as practical challenges in search-based testing. With the recent emergence of large language models, many of these challenges are now thought to be solved. But are they? In this talk, I will explore different applications of search-based testing, together with open, solved, and new challenges.

Dr. Louise Travé-Massuyès
LAAS-CNRS

Louise Travé-Massuyès holds the position of Directrice de Recherche at CNRS, affiliated to LAAS-CNRS, Toulouse, France. She graduated in control engineering in 1982 and received her Ph.D. degree in 1984, both from INSA, Toulouse, France. Her research interests are all related to diagnosis reasoning, tackled by model-based and data-driven approaches. This theme, which she developed throughout her career, led her to consider various formalisms, from both the Artificial Intelligence and the Automatic Control fields, to address problems related to diagnosis. She has been responsible for several industrial and European projects and published more than 250 papers in scientific journals and conference proceedings and 4 books. She is a member of the International Federation of Automatic Control (IFAC) Safeprocess Technical Committee. She holds the chair "Synergistic transformations in model-based and data-based diagnosis" in the Artificial and Natural Intelligence Toulouse Institute (ANITI), France, and serves as Associate Editor for the well-known Artificial Intelligence Journal.

Title: Leveraging the properties of the Christoffel function for anomaly detection in data streams

Abstract: The Christoffel-Darboux Kernel and the associated Christoffel function (CF) are well-known tools from the theory of approximation and orthogonal polynomials. Although they have been largely ignored in discrete data analysis, recent results show that they have many potential uses in data analysis, including applications in machine learning. Some peculiar properties of the CF can be leveraged for anomaly detection. Anomalies, also defined as outliers or out-of-distribution observations, are essential to be detected in data as they can indicate data corruption or faulty behavior. Trust in Artificial Intelligence systems depends on this because their reliability relies on inputs lying in the training distribution. On the other hand, anomaly detection plays a crucial role in certifying data obtained from sensors or images, as well as in identifying symptoms that can be used to drive diagnosis reasoning and health management. This talk presents two methods devised for anomaly detection in streaming data. The first one is DyCF (Dynamic Christoffel Function method), which benefits from incrementality and the ability to deal with concept drift, i.e., updating the model so that it adapts to the distribution. The second method, called DyCG (Dynamic Christoffel Growth method), leverages convergence properties of the Christoffel function so that it is downright tuning-free. Those two methods benefit from a clean algebraic framework and nicely fulfill the data stream requirements related to the non-stationarity of the distributions and infinitely growing data. An evaluation against state-of-the-art methods using synthetic and real industrial datasets shows that DyCF and DyCG outperform not finetuned methods and are clearly better with respect to execution time and memory use.

Dr. Mihai Nica
AVL List GmbH

Dr. Mihai Nica is the Global Head of ADAS/Automated Driving & Connectivity at AVL List GmbH, Graz, Austria. He has more than 17 years of work experience in the field of verification and validation for cyber-physical systems, ten years of experience in the area of EU R&D project management, and more than six years of experience in the area of reliability engineering and functional safety development and management. He worked on over 60 customer R&D research projects with partners from over 30 countries, such as China, India, and Japan. Dr. Nica received his Ph.D. in software debugging and testing from Graz University of Technology.

Title: Navigating the Road to AV Reliability: A Comprehensive Look at Testing Strategies in ADAS and Automated Driving Systems

Abstract: Nowadays, vehicles are becoming more complex than ever, and different powertrain technologies (more complex but efficient engines, hybrid technologies, and electrification) are operated by highly complex software systems using highly efficient power electronics and sensors. Furthermore, the vehicle becomes more and more an entertainment system on wheels, where advanced hardware technologies and software need to coexist robustly against failures and offer reliable experience with respect to intended functionalities. Things get even more challenging when advanced driver assistance functions, connectivity inside and outside the car, or even autonomous driving functions are part of the vehicle-provided functionalities. To address the reliability challenges of such advanced cyber-physical systems, AVL, as a trusted partner for OEMs and TIERs worldwide, offers advanced methods for both developing reliable systems and investigating and testing such cyber-physical systems. This presentation will give insights into such approaches that make use of ODD information, AI, and combinatorial-based testing to enable a comprehensive test program with the highest coverage of the parameter.

Dr. Juha Tiihonen
Variantum

D.Sc. (Tech.) Juha Tiihonen is Lead Developer of sales configuration systems at Variantum oy, Espoo, Finland. He has a strong research background at the University of Helsinki, Aalto University, and Helsinki University of Technology. His interests include product and service configuration, including modeling, configurators, and business processes; software product variability and its management, recommendation technologies, and mass customization. His strengths include conceptual modeling and tool support for the configuration of physical products and services. Juha is a co-founder of Variantum oy, a company specializing in product life-cycle management of configurable offerings.

Title: Solutions for configurable products - from research to VariSuite

Abstract: The presentation outlines a case where research has led to the establishment of a company that provides solutions for configurable products. The VariPDM platform object model and sales configuration modeling conceptualization originated from research at Helsinki University of Technology (now Aalto University). Core configuration modeling concepts fit real products, are understandable, and can be given formal semantics. Parallelly, Austrian research has addressed similar topics. Our cooperation and interaction have benefited us both. For example, we apply the FastDiag algorithm as the backbone of configuration repair functionality, supporting users in case of conflicting requirements. We congratulate TU Graz IST for the first 20 years of academically valued and simultaneously practically applicable research.

Chapter 3

The IST History

1993 – 2001. The Institute of Software Technology has its origins in the call of Peter Lucas to a full professorship for software technology at TU Graz in October 1993. Prof. Peter Lucas came from the IBM research center in San Jose, California, where he worked in the group of John Backus (known for the Backus-Naur-Form) on functional programming languages. Peter was an expert in compiler design and verification but soon broadened his research scope to program development in general. In the 1970s, Peter co-developed the Vienna Development Method, a rigorous method to develop software iteratively from high-level specifications down to code. Soon after his call, in 1994, Peter Lucas became the chair of Formal Methods Europe, an organization promoting program design based on formal logic, as well as a corresponding member of the Austrian Academy of Sciences. With his strong foundational background, he founded the *Ordinariate of Software Technology*, a sub-unit of Prof. Hermann Maurer's *Institute of Interactive Systems (IICM)* (today, the *Institute of Interactive Systems and Data Science*). Prof. Volkmar Haase and Richard Messnarz from IICM joined the Ordinariate. His first assistants and Ph.D. students were Brigitte Fröhlich, Bernhard Aichernig, and Andreas Kerschbaumer, who also organized a colloquium in his honor [13] before he retired in July 2001. During these years, Peter Lucas and his group had a big influence in the education of computer science at TU Graz: they supervised 6 PhD and 70 MSc students and lectured large courses, like Software Paradigms, Software Technology, Compiler Construction and, most prominently, the first year's courses on programming, where the students started with functional programming (in SML) and evolved systematically towards object-oriented programming (in Java). The tradition of training students in declarative languages, like Scala and Prolog, is still kept alive, and Bernhard Aichernig's research group on Formal Development Methods originates from this period. The offices of the Ordinariate were in Münzgrabenstraße 11, next to the Campus Neue Technik. In 2000, the group moved to its current location at Inffeldgasse 16b. After his retirement in 2001, Peter Lucas moved back to California, where he died on the 2nd of February

2015 with 80 years.

2002 – 2023. With the call of Franz Wotawa as the (inofficial) successor of Peter Lucas in 2002[1], the topic of Artificial Intelligence entered the scene. Interestingly, Franz Wotawa has always worked on the borderline between Artificial Intelligence and Software Engineering. His Ph.D. thesis and Habilitation thesis use methodologies from Artificial Intelligence, particularly Model-based Reasoning, to improve automated fault localization. When Franz Wotawa arrived, Andreas Kerschbaumer and Bernhard Aichernig were still available and organizing the Ordinariate exceptionally well. The open position was filled by Roderick Bloem in 2002. Roderick habilitated at the Institute of Software Technology in 2008 but moved to another institute as a full professor in close, timely proximity.

The time between approximately 2000 and 2005 was dominated by the establishment of a new university law (UG 2002) bringing autonomy to universities[2]. Because of these changes, Hermann Maurer and Franz Wotawa decided to split the institute and found a new one, i.e., the Institute of Software Technology in 2003 with a core team of 1 professor, 3 assistant professors, 1 secretary, and 1 computer administrator. However, there was still one open position for a full professor, which was filled in 2003 by Wolfgang Slany, who also negotiated for another assistant professor. The institute started to grow. However, the institute's growth was never driven by positions paid directly by the university. When Franz Wotawa was hired, he also brought two FWF projects with him. One dealt with Java debugging, and the other with fault localization for hardware design languages, particularly VHDL. All these projects come with additional Ph.D. positions. Between the founding year of the institute in 2003 and the year 2023, the institute carried out more than 140 projects funded by different sources, from the European Union, FFG, and FWF to the Christian Doppler Laboratory Association. All institute members have been very active in providing successful project applications, which finally led to the size of the institute being the research home of 63 people in 2024.

But let us return to the institute's history, which Franz Wotawa led until 2010. Between 2002 and 2006, Andreas Kerschbaumer left, and Bernhard Aichernig left for 4 years to fill a position at the United Nations University in Macao. Open positions were filled with Oswin Aicherholzer and Gerald Steinbauer-Wagner. Oswin was originally with another institute and was recommended. Gerald Steinbauer-Wagner was hired to extend the institute's portfolio to autonomous mobile robots. In 2006, Oswin Aichholzer became an associate professor after his successful habilitation. Gerald Steinbauer-Wagner and Franz Wotawa

[1] Officially, Franz Wotawa is the successor of Volkmar Haase. The first call for a successor of Peter Lucas was not successful, leading to a second call after Franz Wotawa already arrived at TU Graz on December 3rd, 2001.

[2] Before the UG 2002, the institutes were able to make external contracts as separate units. Afterward, it is the university directly to have external contracts. Institutes make contracts on behalf of the university. There are, of course, more substantial changes coming with the UG 2002.

organized the RoboCup event in Graz in 2009, which is a big event bringing researchers and school kids working on different kinds of robots together in one place every year. The event was huge, with a budget of about 1 million Euros and about 2000 participants. In 2009, a new full professor, Alexander Felfernig, was hired, also coming with another assistant professor to our institute. From 2010 to 2015, Franz Wotawa was the Dean of the Faculty of Computer Science, which became the Faculty of Computer Science and Biomedical Engineering via successfully integrating institutes from biomedical engineering in the same Dean's period. Therefore, Wolfgang Slany became the head of the institute from 2010 to 2019. During this time, Bernhard Aichernig and Gerald Steinbauer-Wagner habilitated in 2012 and 2017 respectively. Moreover, Franz Wotawa organized an open call for a female assistant/associate professorship funded partially by the Faculty of Computer Science and Biomedical Engineering and NAWI Graz with the goal of installing computer science teaching for NAWI Graz. Birgit Vogtenhuber was hired, and she decided to join our institute. In 2019, Birgit Vogtenhuber finalized her habilitation and became an associate professor. It is worth noting that another initiative of Franz Wotawa when being Dean led to another full professorship position. There was a close collaboration with the University of Klagenfurt and the Dean Gerhard Friedrich. As a result, a project was established to interchange lectures and finally to share a professor position. From 2018 to 2023, Martin Gebser filled this position and brought more resources to our robotics group. Unfortunately it was not possible to keep the position and the group after the project terminated. In 2019, Franz Wotawa came back as the head of the institute and still fills this position, aiming to keep the institute at the forefront of research and provide excellent output in computer science research and teaching. Last but not least, it is worth mentioning that in 2020 and 2021, there were two open calls for a tenure-track position. The two successful candidates Johannes P. Wallner and Yannic Maus, joined the institute in the year 2021. Both already submitted their habilitation theses, and became associate professors in 2024.

Important project activities (2003-2023) We do not want to discuss all projects. However, some are exceptional and worth describing in more detail. The first one is Softnet Austria funded by the FFG. Softnet Austria was within the COMET program of FFG with the objective of bringing software engineering research into practice. Hence, the project covers a wide range of academic partners (TU Graz, TU Vienna, Univ. Vienna, Univ. Innsbruck, SCCH,...) and companies (Siemens, Kapsch, NTT Data,...) coming with an overall budget of almost 16 million Euros. The project allowed the institute to hire several Ph.D. students and increase its growth. The second one is the Christian Doppler Laboratory for Quality Assurance Methodologies for Autonomous Cyber-Physical Systems (QAMCAS) funded by the CDG as the funding agency representing the Austrian Government and AVL as the industrial partner. From 2018 to 2024, QAMCAS has funded 4 Ph.D. students, one Post-Doc, and 1-2 Master's students working on smaller projects. QAMCAS has also shown that providing foundational research is very beneficial for companies.

Scientific events carried out (2003-2023)

DX 2002	The 13th International Workshop on Principles of Diagnosis (DX) was carried out at the Hotel Panhans, Semmering with about 40 researchers working on diagnosis.
QR 2005	The International Workshop on Qualitative Reasoning (QR) was organized at TU Graz with more than 30 participants.
EuroCG 2007	The 23rd European Workshop on Computational Geometry was organized at TU-Graz with over 100 participants from 19 different countries worldwide. The welcome reception took place in the city hall, hosted by the mayor.
RoboCup 2009	RoboCup founded in 1997, the largest and most important event for intelligent and autonomous robots, took place June/July 2009 in Austria for the first time. It attracted about 2,500 participants from more than 40 countries all over the world. The event was carried out at Messe Graz and was a big success.
DX 2014	The 25th International Workshop on Principles of Diagnosis (DX) was located in Graz. Again more than 40 researchers meet to discuss recent challenges and solutions of model-based diagnosis.
ICST 2015	The 8th IEEE International Conference on Software Testing, Verification and Validation with about 250 participants from all over the world.
ICTSS 2016	The 28th International Conference on Testing Software and Systems. This event was carried out in October 2016 with about 70 participants.

ConfWS 2018	20th International Workshop on Configuration with around 40 participants.
IEA/AIE 2019	32nd International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems. This three days event was carried out at the main building of TU Graz and had about 100 participants.
ISMIS 2020	25th International Symposium on Methodologies for Intelligent Systems, Sep. 23rd-25th 2020, conference with around 80 participants (online event due to Covid-19).
ConfWS 2022	24th International Workshop on Configuration with around 40 participants.
SPLC 2022	26th ACM International Systems and Software Product Line Conference, September 12-16, 2022. An international conference with around 170 participants.

Spin-Offs (2003-2023)

DiLT Analytics GmbH	founded 2021	DiLT Analytics work on smart software solutions for efficient buildings and energy systems. `www.dilt.at`
SelectionArts GmbH	founded 2013	Company with a focus on developing recommender systems based e-learning solutions. `www.selectionarts.com`

Guest in our Seminar Room (2003-2023)
There have been a lot of talks given by our colleagues from all over the world. In the following, we only mention a few.

Dines Björner, Prof. emeritus TUD, Co-founder of formal methods, December 1st, 2008.

John Hughes, Prof. Univ. Chalmers, Property-based testing, , April 23rd, 2013.

Nina Yevtushenko, Prof. Univ. Tomsk, Model-based testing, February 2nd, 2013.

Manuel Nunez, Prof. Univ. Complutense Madrid, Model-based testing, June 11th, 2013.

Alexander Kleiner, Bosch, Mobile Robotics, RoboCup, July 2nd, 2013.

Radu Madescu, INRIA, Model-based testing, March 3rd, 2016.

Eric Wong, Prof. Univ. Texas, IEEE Reliability Society, debugging, combinatorial testing expert, May 3rd, 2016.

Dimitris Simos (on the right), SBA Research, combinatorial testing and security testing expert, June 20th, 2017.

Johann De Kleer, Parc, Co-Founder of the model-based and qualitative reasoning community, July 11th, 2018.

Chapter 4

Research at IST

The Institute of Software Technology comprises ten research groups. In this chapter, each group outlines its structure and research activities. For the research interests of the group leaders have a look at Figure 4.1. Many members of the institute provide research and teaching activities in more than one area of research.

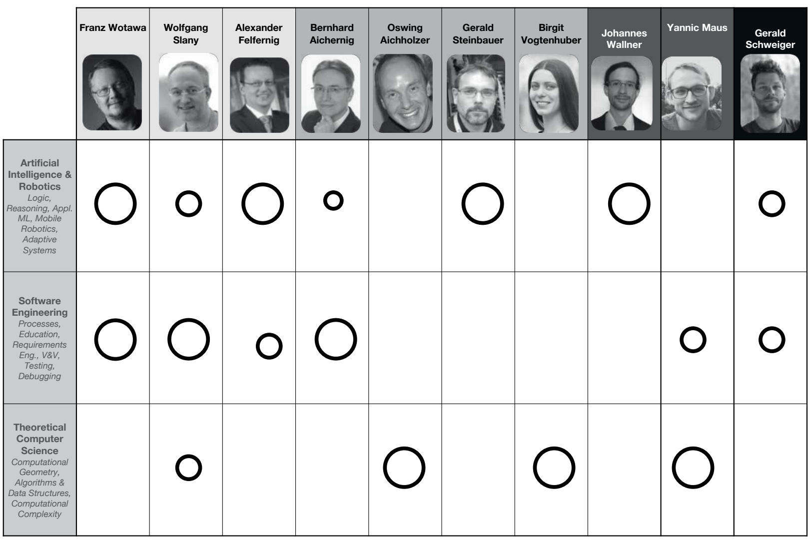

Figure 4.1: Research group leaders and their research topics.

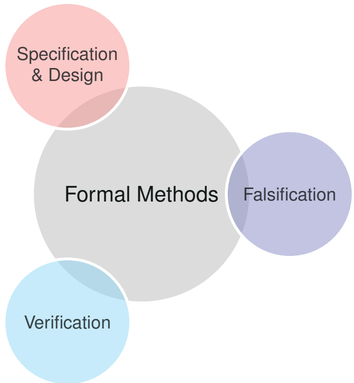

Figure 4.2: Research area of the Group Aichernig.

4.1 Group Aichernig - Formal Development Methods

4.1.1 Research Focus

The group of Bernhard K. Aichernig focuses its research on formal methods for software development. Formal methods are mathematically rigorous techniques for the specification, development, analysis, and verification of computer-based systems. The group's work covers activities in all development phases, ranging from requirements engineering, over testing and verification, to reverse engineering.

As shown in Figure 4.2, the research activities can be broadly grouped into three categories: (1) *Specification & Design* concentrates on the early phases of development and considers objectives, such as finding a suitable abstraction for a given problem. *Falsification* mainly deals with software testing. Research belonging to this category offers for instance methods for automated test-case generation, with the goal of efficiently finding defects. *Verification* in contrast aims to show that a software system or component does not contain flaws. For example, proof-based development for instance provides techniques to create software that is correct by design.

The main focus of the group was in the past model-based testing and automated test-case generation. In recent years the group has researched combinations of Formal Methods and Artificial Intelligence, including learning-based testing, model-learning for re-engineering, and the testing and analysis of machine-learned components. We use our extended competence in model-based test-case generation in order to advance the state of the art in active automata learning, a form of model learning where data is gathered from a system during the learning of its behavioral model.

4.1.2 Group Structure

At the time of writing, the group consists of Bernhard K. Aichernig (associate professor), Martin Tappler (postdoc), and the PhD students Edi Muskardin, Benjamin von Berg, and Felix Wallner. Andrea Pferscher completed her PhD studies a few months ago and is now a postdoc at the University of Oslo. The research is financed by three ongoing projects:

- DES Lab: The TU Graz - SAL Dependable Embedded Systems Lab is a collaboration with Silicon Austria Labs. Together with other colleagues from SAL and TU Graz, we carry out foundational research on the intersections of formal methods and AI. Bernhard is the co-leader of the lab and Martin Tappler and Edi Muskardin are financed through this cooperation.

- LearnTwins: This national project on Learning Digital Twins for the Validation and Verification of Dependable Cyber-Physical Systems is an applied research collaboration with AVL and AIT. In this project, we combine deep learning with automata learning in order to automatically synthesize digital twins of AVL's measurement devices. Felix Wallner is working on this topic.

- AIDoArT: In this EU project on AI-augmented automation supporting modeling, coding, testing, monitoring, and continuous development in Cyber-Physical Systems we cooperate with AVL on two use-cases: (1) we apply automata learning for inferring a model of human drivers for more realistic simulations of driving emissions (real driving emissions); (2) we combine search-based and learning-based testing for testing autonomous vehicles. Benjamin von Berg is working on these topics.

With respect to teaching, our group covers three large bachelor courses: Software Paradigms, Declarative Programming, and Quality Assurance in Software Development. In the study year 2022/23, these courses had between 150 and 300 students enrolled. In addition, we offer smaller master courses on Model-based Testing and Formal Specification and Design of Software.

4.1.3 Scientific Highlights

Historically, the group evolved from formal modeling with specification languages, over model-based testing to its current research focus on model learning. In the following, we highlight the two most prominent techniques that originated in the group:

Model-based mutation testing. Mutation testing is a way of assessing and improving a test suite by checking if its test cases can detect a number of injected faults in a program. The faults are introduced by syntactically changing the source code of a program. It is considered the golden standard when it comes to test coverage.

In our work, we generalized mutation testing from program testing to model-based testing. Like in model-based testing, the aim is to use the model for both, generating test cases and as a test oracle. Hence, we automatically generate test cases from a test model and test if a system under test conforms to it. In contrast to classical model-based testing, only those test cases are generated that find the injected faults in a set of mutated models. The generated tests are then executed on the system under test and will detect if a mutated model has been implemented. Hence, model-based mutation testing rather tests against non-conformance than for conformance. In terms of epistemology, we are rather aiming for falsification than for verification. It is a complementary fault-centered testing approach.

Our work has led to a series of research results that range from theory via implementation to application. The theoretical investigations are based on formal semantics of models, conformance, and test cases. We showed that the search for mutation test cases can be generalized to a non-conformance problem. The result is a theory of mutation testing that also applies to non-deterministic models. We have presented this theory in the semantic styles of Back's Refinement Calculus [7] and in Hoare and He's Unifying Theories of Programming [14].

We implemented these theoretical results for different modeling styles: relational specifications [20, 15], process algebras [195], coordination languages [145], Action Systems [16, 12], Qualitative Reasoning models [11], and Timed Automata [17]. The underlying techniques for implementing our test case generators for this variety of models includes constraint solvers, SMT solvers, bi-simulation checkers, symbolic execution, concolic execution, term rewriting, input-output conformance checkers, and Qualitative Reasoning simulators.

Learning-based Testing. In this line of research, we exploit automata learning to detect failures in reactive systems, including machine-trained controllers. The idea is to learn a behavioral model of the system-under-test from interactions with the system. This model can then be verified in order to detect failures. In passive learning, we learn from given system logs. In active automata learning, we query the system until we cannot find mismatches in the behavior of the model and the system. The latter is an iterative process where model learning and model-based testing are interleaved. Hence, the group can build upon its expertise in model-based testing.

This research started with an experiment of testing MQTT servers [187]. As a result, we could find several bugs in existing open-source servers. This initial work spanned a variety of research directions: Inspired by mutation testing, we contributed an improved test-case generation technique for active automata learning [21].

Then, we investigated the learning of richer modeling semantics: In order to overcome the limitation of deterministic systems, we developed the first active automata learning algorithm for stochastic systems [186]. We also showed how model learning supports the synthesis of controllers by combining it with probabilistic model checking [22]. For timed systems, we developed a passive and active genetic algorithm to learn timed automata [188, 19]. In order to test and predict the response time of systems, we combined model-based testing with the learning of response-time distributions [9]. This was our first work on integrating classical machine learning into our testing techniques.

Machine learning soon became the center of our attention: We showed that we can use automata learning in order to generate better training data for deep learning [10]. We won the best paper award. Then, we developed a novel testing method for agents trained with reinforcement learning [189] — our first IJCAI paper. Most recently, we extended this work to improve an agent with the help of generated test cases [185] — one of our two first ICSE papers. The other one being on learning models from noisy data [192].

Finally, we developed a new security testing method based on automata learning [18, 164]. This work uncovered several vulnerabilities in MQTT servers and Bluetooth Low Energy devices.

4.1.4 Impact Highlights

Over the years, our work on automated testing resulted in the discovery of many bugs in real-world applications, including web servers, MQTT brokers, and Bluetooth devices. More important is the general impact, that our techniques aim at reducing the efforts in testing — a process that demands 50% to 70% in critical projects. With model-based testing, we free the tester from writing test cases manually. With learning-based testing, we also free the tester from modeling.

The work on model-based mutation testing led to the MoMuT tools[1] [6], now maintained by our research partner AIT. AVL uses it to generate integration tests for its automotive measurement devices [8].

Our research on learning-based testing led to the development of AALpy[2], our own library for automata learning implemented in Python [150]. With its support of learning deterministic, non-deterministic, and stochastic models it receives growing attention in the international research community.

The recent publications at IJCAI [189] and ICSE [185, 192] show that our group is on the right track in bringing our background from formal methods into Artificial Intelligence and Software Engineering.

[1] `https://momut.org`
[2] `https://github.com/DES-Lab/AALpy`

4.2 Group Aichholzer - Discrete and Computational Geometry

4.2.1 Research Focus

The research interests of the group of Aichholzer include discrete and computational geometry, data structures and algorithms, combinatorial properties of geometric and topological graphs, enumeration algorithms, and combinatorial games.

On the algorithmic side, we are especially interested in combinatorial properties of triangulations and related data structures to obtain efficient algorithms for transforming or counting/enumerating triangulations. This includes consideration of additional restrictions like bounds on the maximum face or vertex degrees. In the area of discrete geometry, we consider typical Erdős-type problems on empty and non-empty convex polygons spanned by (colored and uncolored) point sets in the plane.

On (geometric) graphs the focus lies on reconfiguration problems and the minimum crossing number of complete graphs. Investigating the differences between geometric graphs (vertices are points in the Euclidean plane and edges are segments connecting two points) and general (topological) drawings of graphs has been proven very fruitful in the last few years and will also play a central role in the next years. Moreover, the relation of crossing minimal drawings to combinatorial structures, like k-edges and order types in the geometric case, or rotation systems in the topological setting, are of special interest.

On the teaching side, our group covers algorithms, data structures, discrete geometry, and combinatorial aspects of enumeration and games.

4.2.2 Group Structure

The group consists of Oswin Aichholzer (associate professor) and a fluctuating number of PhD students employed by funded projects. There is no further support for permanent scientific positions by the university or faculty for this group. During the last years, 10 PhD students obtained their PhD in our group, and most noteworthy already three of them got permanent tenure track positions at different universities.

Within TU-Graz our group has strong connections to other faculties, especially mathematics and electrical engineering: Aichholzer was Dean of studies of a joint curriculum between computer science and electrical engineering (telematics) and is in the leading team of the FoE ICC (Fields of Expertise Information, Communication, and Computing) which combines the three faculties computer science, mathematics, and electrical engineering. Moreover, Aichholzer is a member of the Doctoral School of Mathematics, and founding member of the Graz School of Discrete Mathematics (GSDM), and PI of joint projects under this umbrella. See Section4.2.4 for more details on those projects and also the connection to the international community via collaborative projects.

4.2.3 Scientific Highlights

In the last 20 years, the main focus of our group has been on combinatorial structures behind point sets ((abstract) order types) and simple topological drawings of graphs ((abstract) rotation systems) and their relation to each other. Some exemplary results are given in the next paragraphs:

Order types. Many problems in discrete geometry are based on straight line graphs where the vertices are points embedded in the plane and edges are straight line segments connecting them. All properties of such graphs are determined by the position of the points. If we are interested only in combinatorial properties of the graph, like if two given segments intersect or if a cycle in the graph forms a convex polygon, then this information can be read of the triple orientations of the points. The set of all these triple orientations is called the order type of the point set and is also known in the context of rank three oriented matroids. For abstract order types (a collection of triple orientations that fulfills certain criteria) it is $\exists\mathbb{R}$-hard to decide if it is realizable as a set of points in the plane. This makes it challenging to apply the vast theory that exists for abstract order types to geometric graph problems. One of our contributions in this area is the generation of all realizable order types for point sets of small cardinality, and their application to different problems [24, 28]. For example, we were able to significantly improve the best straight line drawings minimizing the rectilinear crossing number [31] (see also below), improve bounds for the existence of convex k-gons in point sets [25], or get improved results for crossing families [33]. Moreover, the sets proved useful to generate counterexamples to several open research conjectures in combinatorial geometry or to build an induction base for asymptotic improvements. An example of the latter is improved lower bounds for the number of triangulations of a set of points in general position [27]. The data has been made available to the community, and still, new applications and results are discovered and I get requests for special settings which are close to order types.

Rotation systems. Graphs drawn as simple topological graphs (now edges do not need to be straight lines anymore but can intersect at most once, either in a common endpoint or in a proper crossing in their interior) are a common way to visualize graphs. In contrast to straight-line graphs (see previous paragraph) such graphs are very hard to handle with a computer (how can one store lines, that almost arbitrarily meander in the plane, in an efficient way?). For complete graphs and if we are solely interested in combinatorial properties of the graph (for example, and similar to before, if we want to know if two given lines have a common crossing) the so-called rotation system (the collection of the order of emanating edges for all vertices) turned out to be useful. Generating all realizable rotation systems for graphs of some (limited) constant size turned again out to be a source of helpful information [2]. Results for the crossing number of the complete graph can be derived from this data [1] (see next paragraph), but also answers to many combinatorial problems on simple topological graphs can be derived [32]. One key is that the exhaustive knowledge of rotation systems and their drawings of small cardinality often allows us to observe the underlying structure and then generalize this to the unbounded case.

Crossing numbers. Minimizing the number of crossings in a drawing of a graph is a central topic in graph theory. For complete and complete bipartite graphs there are long-standing open problems and over 60-year-old conjectures (most prominent the Harary-Hill Conjecture or Zarankiewicz's Conjecture) on the minimal possible number of crossings and the structure behind the drawings obtaining them. These questions concern all types of drawings, where our research concentrated on geometric graphs and simple topological graphs. Using the results from the two previous paragraphs we have been able to make various contributions. This includes a series of papers showing that the 60-year old Harrary-Hill conjecture holds for many graph classes, thereby introducing the concept of shellable and bi-shellable drawings [1, 4, 3]. We also showed that geometric crossing minimal drawings have a triangular convex hull [26], and give the best asymptotic upper bounds on their crossing number via the computation of large examples [31].

Reconfiguration of graphs. Transforming a graph via steps of bounded complexity into another graph of the same graph class is called reconfiguration. A typical reconfiguration operation is exchanging one edge of the graph for another edge, which is often called a flip. The flip graph is defined as the graph having a vertex for each configuration and an edge for each flip. Three questions are central: studying the connectivity of the flip graph, its diameter, and the complexity of finding the shortest flip sequence between two given configurations. Many of the reconfiguration questions are central to theoretical computer science, like the reconfiguration of trees (e.g. the rotation operation in binary trees which are a common data structure), or the transformation of triangulations when used to visualize surfaces or 3D models [23]. Our contributions range from hardness results (we showed that the problem of optimizing the reconfiguration distance between two triangulations of simple polygons is NP-complete [34]) to improved upper bounds for the reconfiguration of trees, matchings, or token graphs [30, 29].

Combinatorial Games. Combining enumerative combinatorics with efficient data structures and algorithms is one of the core competencies of our group. Using this with results from game theory we derived a method to exhaustively analyze several popular board games. This is not only a very motivating topic for teaching efficient algorithms, but can actually help even professional players improve their skills. If you think now that this is just “yet another program that plays some game better than any human being” let us point out that our technique is superior to the usual AI approaches. The difference is that our results allow the program to play a game in a mathematically provable perfect way. Not only does it provide for any possible game situation a winning move (if one exists), but it can tell for any position and any possible move how good the moves are. More concretely that means that we know how many moves (to be precise, in game theory actually the term half-move would be used, but we refrain from giving too much details here) the program guarantees a win, regardless of how good (or bad - no assumptions on the second player are made) the opponent performs.

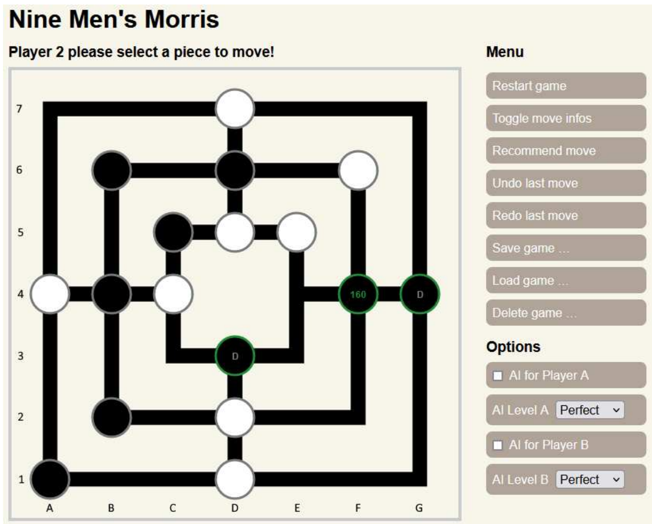

Figure 4.3: A game state of Nine Mens Morris where the second player (black) has a guaranteed win in 160 moves. Black has to continue with F4 - E4.

For example, for the shown position of Nine Mens Morris in Figure 4.3 it is known that the program will win in at most 160 moves. The next move to win is sliding the token from F4 to E4. All other possible moves would not guarantee a win but allow for a draw. The given number of needed moves is optimal in the sense that a perfectly playing opponent can postpone the loss of the game precisely by that number of moves. This shows the power of our approach - just imagine a classical AI-based method to “think” 160 moves ahead in an exact way.

If you are interested in playing against (or learning from) our programs, you can do so at http://ninemensmorris.ist.tugraz.at:8080/ for Nine Men’s Morris (german: Mühle), http://connect4.ist.tugraz.at:8080/ for Connect Four (German: Vier Gewinnt), and http://hexapawn.ist.tugraz.at/ for Hexapawn (a pawn-based chess variant). Have fun!

4.2.4 Impact Highlights

Significant impact has been made via international joint projects and the foundation of several international annual workshops. In the following, we list projects in the last 20 years where at least one additional university (besides TU-Graz) was/is involved (in chronological order).

- 2005 - 2007: PI of subproject ''Computational Geometry" within the National Research Network "Industrial Geometry", funded by FWF with 400k€ for subproject. Collaboration of four Austrian universities (Graz, Vienna, Linz, Innsbruck). In the context of advanced geometric representations the synergies of approximation theory and computational geometry were studied in this subproject. See http://industrial-geometry.ist.tugraz.at/project05.php
- 2007 - 2008: Head of a bilateral grant France-Austria "Amadée" funded by Österreichischer Austauschdienst "Konzepte der Rechnerischen Geometrie - Theorie und Anwendungen". Cooperation with INRIA Sophie Antipolis, France on several topics from Discrete and Computational Geometry.
- 2008 - 2009: Head of a bilateral grant Spain-Austria "Acciones Integradas" funded by Österreichischer Austauschdienst "Optimierung geometrischer Graphen". Cooperation with Universidad de Alcala, Spain on concepts of geometric graphs and crossing numbers.
- 2008 - 2011: Head of National Research Network (aka SFB Special Research Programme) "Industrial Geometry". Funded by FWF with 2.39M€. Collaboration of four Austrian universities (Graz, Vienna, Linz, Innsbruck). Industrial geometry is based on computational techniques which originated in various areas of applied geometry. This project combined experts from areas like Computer Aided Geometric Design, Computer Vision, Image processing, Computational Geometry and Robot kinematics. See http://industrial-geometry.ist.tugraz.at.
- 2011 - 2014: Head of the multinational Collaborative Research Project "ComPoSe". Supported by the EUROCORES program EuroGIGA of the European Science Foundation (ESF) with 1.57M€. Collaboration of several European universities (TU Graz, Austria; Université Libre de Bruxelles (ULB), Belgium; Technische Universität Berlin, Germany; Universitat Politècnica de Catalunya; Alfréd Rényi Institute of Mathematics, Hungary, École Polytechnique Fédérale de Lausanne; Charles University, Prague, Czech Republic; ETH Zürich, Switzerland). ComPoSe focuses on combinatorial properties of discrete sets of points and other simple geometric objects primarily in the plane. See http://www.eurogiga-compose.eu
- 2017 - 2022: Work package leader and PI for CONNECT, European Union's Horizon 2020 research and innovation program, Marie Skłodowska-

Curie grant. Leader of Work Package 1: Geometric Networks. Funding 100K€ for WP 1. This is a collaboration of 14 universities in Europe and North- and South America supported by the EU. This project aims to obtain new insights into the behavior of networks, which are studied from a geometric and computational perspective. Thereto, the project brings together researchers from different areas such as Computational Geometry, Discrete Mathematics, Graph Drawing, and Probability. See https://www.connect-rise.eu

- 2018: Head of the Austrian Team of the Japan-Austria Bilateral Seminar: Computational Geometry Seminar with Applications to Sensor Networks supported by the Austrian Science Fund (FWF) and the Japan Society for the Promotion of Science (JSPS) This project is about the theory behind and the application of different types of sensors in various networks. See http://www.dais.is.tohoku.ac.jp/∼jinhee/JA2018.html

- 2015 - 2024: PI of Subproject "(Geometric) graphs: Flip distances and crossing numbers" within the Doctoral Program "Discrete Mathematics" approx 420 000€ for subproject, funded by FWF and supported by TU-Graz. In this doctoral program researchers from TU-Graz, KFU Graz, and the University of Leoben collaborated on different areas of Discrete Mathematics. The main focus was on educating PhD students in this area. The outcome was a total of 70 PhDs and over 300 publications by 17 subprojects. See https://www.math.tugraz.at/discrete/.

- 2024 - 2028: An FWF-funded doc-funds project in cooperation with other members of IST (Birgit Vogtenhuber and Yannic Maus) and the faculties of mathematics at TU-Graz and KFU-Graz. The project will start in 2024 and support 11 PhD students in an collaborative way between the involved principal investigators. The central topic is discrete mathematics and covers diverse fields of mathematics and theoretical computer science. Funded by FWF with approx 2.2M€ for all subprojects.

The organizations of international conferences and the foundation (and also organization) of international annual workshops have a main impact on the research community. As one example let us mention here the foundation of the annual European Research Week on geometric graphs. In 2004 (so shortly after the IST was founded) during the visit of five colleagues to Graz, we started a series of annual workshops on pseudo-triangulations and (later) geometric graphs in general. This is meanwhile established as the European Research Week on Geometric Graphs and takes place annually at different locations in Europe, see, e.g., https://ggweek2023.web.uah.es/ for a nice overview. Participation is by invitation only, and in the latest editions about 30 researchers joined. Among them are several young PhDs from different places all over Europe, which not only shows that the area is still vivid after almost 20 years, but also makes sure that the interest in these research weeks, and thus this research area, will continue.

Aichholzer is also a member of the steering committee for the annual Crossing Numbers Workshop, which is one of the central workshops for the crossing number community. Every year it is organized by a different research group and attracts researchers and young PhDs from all over the world for intensive joint research (http://www.crossingnumbers.org/events.php).

4.3 Group Felfernig - Applied Software Engineering & AI

4.3.1 ASE Background

The research group *Applied Software Engineering & Artificial Intelligence (ASE)* has been established by *Alexander Felfernig* as part of the Institute of Software Technology in March 2008 – since then, the ASE team has successfully completed numerous **European Union and national (Austrian) basic research projects**. If you are interested in a detailed overview of the ASE research output, memberships in scientific committees, and teaching activities, take a look at: **ase.ist.tugraz.at**. The research topics of recommender systems [118], knowledge-based configuration [85], and model-based diagnosis [81] were already core topics of Alexander Felfernig and colleagues in the research group of Gerhard Friedrich at the University of Klagenfurt. We want to emphasize the role model effect of Gerhard Friedrich in establishing a highly competitive and productive research spirit fostering high-quality research projects and publications. We want to thank all of our students for their valuable contributions: with December 2023, **34 ASE students** have successfully completed their master thesis, and **10 PhD students** have successfully completed their PhD project.

We want to mention persons who – over the years – contributed to the ASE research landscape as it is now: first of all, *Carmen Riedler* started (and successfully completed) a master thesis on the application and adaptation of knowledge-based configuration technologies in financial services. The outcomes of her thesis can be regarded as a major input for establishing the research field of constraint-based recommender systems [79, 82] which is still a highly relevant research area especially when dealing with complex products and services. On the basis of these insights combined with the results of a European Union research project, **ConfigWorks** has been founded with *Gerhard Friedrich*, *Alexander Felfernig*, *Markus Zanker*, *Christian Russ*, and *Dietmar Jannach* in 2003 with a focus on the development of constraint-based recommender solutions for complex products and services. In this context, we want to thank *Dieter Wedenig* and *Andrea Knapp-Rebernig* for their valuable industrial input that helped to advance Configworks software solutions in terms of feature relevancy. Also thanks to *Bartosz Gula* and *Gerhard Leitner* from the University of Klagenfurt for their research contributions to bring together the fields of machine learning and cognitive psychology and to *Juha Tiihonen* for the long-lasting and fruitful cooperation with the Aalto University and the University of Helsinki. Special thanks go to *Walter Svoboda* as a managing director of *Wüstenrot Datenservice* who was a highly cooperative and supportive industrial partner with highly relevant and innovative research ideas leading to numerous national and international research projects. Core topics of the cooperation with Wüstenrot were recommendation services for integrated financial planning and human-aware recommender interfaces taking into account knowledge and theories of human decision-making in AI systems.

4.3.2 ASE Research & Achievements

The mentioned research fields of recommender systems, knowledge-based configuration, and model-based diagnosis are still a central focus of ASE research. Since we are simply not able to discuss all the ASE research contributions – by the end of 2023, the **ASE research publication count is over 400** including the books on "Recommender Systems" (Cambridge University Press), "Group Recommender Systems" (Springer), and "Knowledge-based Configuration" (Morgan Kaufmann) – we limit ourselves to the mentioning of developed *new fields of research* (see Figure 4.4). Major developments are **(a)** the integration of symbolic and sub-symbolic AI, for example, in terms of directly *integrating recommender systems (recommendation models) with constraint-based reasoning* [166], **(b)** the development of *direct and parallelized model-based diagnosis* approaches alleviating the retrieval of preferred diagnoses [83, 87, 88, 137], **(c)** the development of solutions to extend and integrate *recommender systems with software engineering* [78], for example, in the context of different requirements engineering tasks [92], **(d)** *sustainability-aware recommendation approaches* [91, 191], **(e)** and *psychology-aware recommender systems* able to integrate theories of human decision making into recommendation processes [80, 181, 190]. All related research results have been developed within the scope of basic research projects in the line of *FFG Bridge* and *European Union Horizon* projects (**with an ASE funding of around 4 Mio EUR**) and published in books and internationally renowned conferences and journals (see ase.ist.tugraz.at). In order to exploit these scientific results in industry, **SelectionArts** has been established as a spin-off of the ASE research group in 2013. The company focuses on the development of recommendation-supported learning technologies applied in various Austrian public institutions [182]. The driving force behind this is *Martin Stettinger* (PostDoc researcher) who is an enthusiastic researcher with a heavy interest in developing and also applying research results in industrial settings.

Awards received by members of the ASE research group are the following: the **Kärntner Innovationspreis** (Alexander Felfernig with the co-founders of ConfigWorks, 2003), **Kulturpreis des Landes Kärnten** (Alexander Felfernig, 2008), **Heinz Zemanek Preis of the Austrian Computer Society** (Alexander Felfernig, 2008), and the **Houska Award Nominee** (ASE team, 2014). Besides these, the ASE research group has received **various best paper awards at**

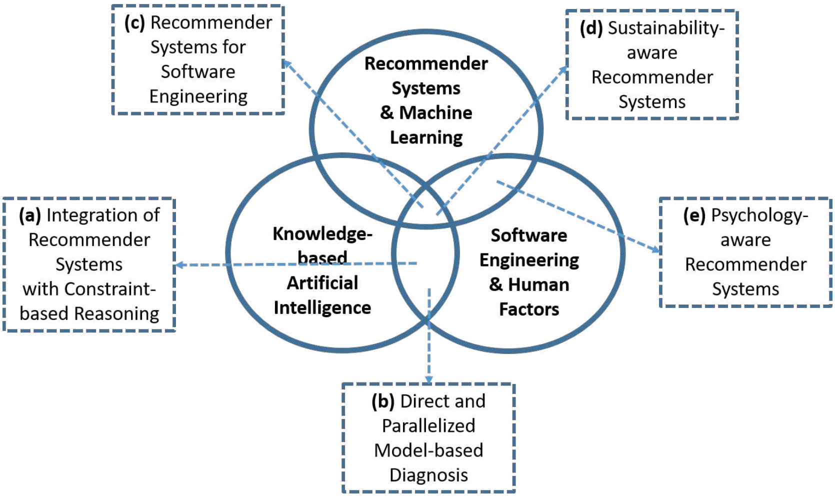

Figure 4.4: Overview of Applied Software Engineering & AI research topics. Related industrial deployments: (a) Constraint-based Financial Service recommendation (e.g., **wuestenrot.at** [86]), (b) Direct diagnosis solutions applied in various industrial settings (e.g., **selectionarts.at** [136, 137]), (c) utility-based and group recommendation for open source development (e.g., **eclipse.org** [89]), (d) group recommenders for open innovation in the public sector (e.g., in the country of Carinthia with **eventhelpr.com** [84]) and recommender systems for technology-enhanced learning (e.g., **knowledgecheckr.com** [182]), and (e) biases-ware group decision making (e.g., **CHOICLA** group decision support [181] and **sports recommender systems** [90]) .

scientific conferences and also organized scientific conferences such as the ACM Conference on Recommender Systems (**ACM RecSys**) and the ACM Systems and Software Product Line Conference (**ACM SPLC**). The **current ASE team (2023)** consists of *Trang Tran* (PostDoc researcher with a strong focus on human decision making, explanations, and recommender systems), *Seda Polat-Erdeniz* (PostDoc researcher with a strong focus on IoT, machine learning, and explanation approaches in the medicine domain, *Martin Stettinger* (PostDoc focusing on the development personalized e-learning technologies), *Viet-Man Le* (PhD student focusing on neurosymbolic computing approaches and the optimization of conflict detection and model-based diagnosis algorithms, for example, on the basis of different machine learning and parallelization approaches), *Sebastian Lubos* (PhD student focusing on the development of new context-aware and LLM-enhanced video recommendation approaches), and *Damian Garber* (PhD student with a focus on configuration space learning with applications in error detection and complex systems parameter optimization

– in cooperation with *Tamim Burgstaller*, who is an ASE student study assistant).

4.4 Group Maus - Algorithms and Complexity

The group performs foundational research in the theory of computation with a strong focus on distributed and decentralized computation. The group has been recently established at Graz University of Technology and is currently growing, including emerging collaborations with other researchers from TU Graz, e.g., from Mathematics.

Theory of computation deals with all kinds of theoretic questions on algorithms, let it be designing faster shortest path or routing algorithms or designing a rigorous complexity theory. Thus, it belongs to the most classic topics of computer science. This is underscored by the widespread inclusion of extensive mandatory coursework on the theory of computation in virtually all computer science university programs around the world.

Our group faces timely challenges in the theory of computation. Many of today's and tomorrow's computer systems distribute data to several machines, or the systems are built on top of large-scale networks such as the Internet or sensor networks. Driven by the necessity to adapt classical research on the theory of computation to our modern world, we build theoretical foundations for these settings. We analyze and develop distributed and decentralized algorithms that are fast and communication efficient, or we rule out the existence of such algorithms. Much of our work is backed up by a thorough complexity theory for such systems.

4.4.1 Research Focus

In order to understand distributed and decentralized computation, we focus on classic graph problems. Graph problems do not only serve as a powerful abstraction for modeling and solving a wide range of real-world problems, but they also help us classify problems into complexity classes, providing insights into computational difficulty. In the more general context, we want to answer the following question:

What makes distributed and decentralized computation hard?

To answer that question, we on the one hand study simple-to-state classic graph problems such as the vertex coloring problem (see below), but also aim to develop a complexity theory that ultimately can be used to classify the complexity of (almost) all local constraint satisfaction problems in distributed and decentralized settings.

Our research has focused on designing or ruling out the existence of efficient algorithms in various standard distributed and decentralized models of computation. Prime examples are the LOCAL and the MPC model. The **LOCAL model** is the standard model for building theoretical foundations of

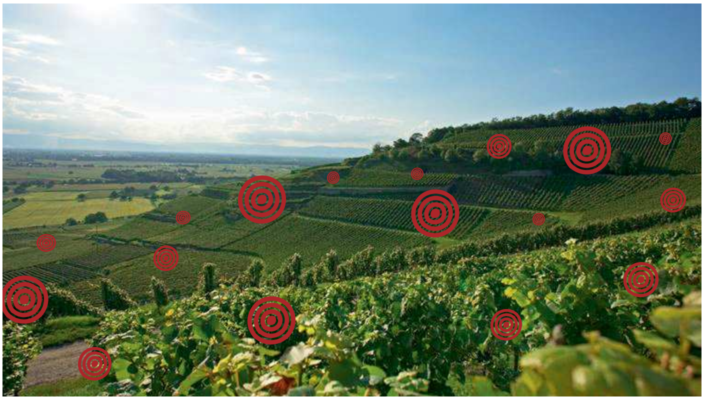

Figure 4.5: Sensor networks are distributed if there is no central base station.

network algorithms [140]. A communication network is abstracted as a n-node graph where vertices are computing entities and edges are communication links. Computers communicate with each other in synchronous rounds, and at the end of the computation, each vertex needs to know its own part of the solution, e.g., its own color in a vertex coloring problem. The crux is, that initially, a node does not know what the network looks like—think of your mobile phone that does not know the topology of the internet. The objective is to design algorithms that require as few rounds as possible, as fewer rounds generally mean less communication.

Originally inspired by decentralized massive data processing frameworks such as MapReduce [63], Spark [209], or Hadoop [198], the **Massively Parallel Computing Model (MPC**) has been introduced to build theoretical foundations for such systems [125]. By now, it has developed into its own active research community, yielding a large presence at top-tier theory conferences and workshops. In the model, many machines, each holding some small part of the input, work together to solve the problem at hand. At first sight, there is a huge discrepancy between the MPC model (all-to-all communication between machines in a cluster) and the LOCAL model (communication with immediate neighbors in the network only). Still, most state-of-the-art graph algorithms in the MPC model are model-optimized implementations of existing LOCAL algorithms. Also, it is known that faster MPC algorithms yield faster LOCAL algorithms.

More generally, research is currently only at an initial stage of forming a rigorous understanding of the connections between different settings where computation is performed with limited information.

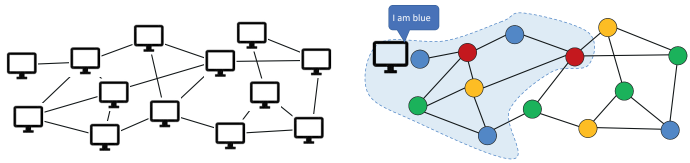

Figure 4.6: Networks are abstracted as graphs.

4.4.2 Research highlights

We focus on a few research highlights of our recent research.[3]

Benchmark problem graph coloring. Graph coloring problems are at the core of the field of limited-information computation and understanding these problems yields ample insights into computational difficulty. Members of the group have co-authored more than ten papers on the topic, presenting various pioneering and state-of-the-art algorithms. For example, we co-authored the first time- and communication-efficient distributed vertex coloring algorithms [45, 103]. To get a feeling for the problem, let us illustrate it with an artificial example. Consider a large network of sensors deployed in the vineyards in the south of Styria that measure parameters such as the temperature or the humidity. Due to the steep nature of the yards, obstacles, and limited battery capacity, these sensors cannot communicate with a central base station. Instead, they build a multi-hop network and a sensor only communicates with the sensors in its vicinity. Further, their communication is wireless, so two nearby sensors cannot use the same frequency to send their messages, as otherwise interference might occur. Thus, our goal is to assign frequencies such that close-by sensors do not receive the same frequency. This can be modeled as a classic *graph coloring problem*: The network is abstracted as a graph where each vertex resembles a sensor and each edge between two vertices indicates that the two sensors are so close that they cannot use the same frequency. The objective is to assign a color (= frequency) to each vertex such that neighboring vertices receive different colors. To save resources, one desires to use a few colors. Yet so simple, graph coloring has played an influential role in computer science, maths, and also in distributed computing. Computing an assignment using a minimal number of colors cannot be solved efficiently unless $P = NP$. In the distributed setting, we use more colors, still the problem remains challenging. Recall our example and how it resembles the LOCAL model. Each vertex of the graph is its own computer and does not know the network. How should one of these decide on a color without a central coordinator? Should it go for blue, red, or green? It has to coordinate with its

[3] Parts of the examples in this section have appeared in a prior text on the research group of Yannic Maus in the magazine *TU Graz People 80/2021-4*.

neighbors to not pick the same color, and its neighbors have to coordinate with their neighbors, etc. Where do we begin? Roll a dice!

Complexity theory & the role of randomness. Let each node simultaneously pick a random candidate color by rolling a dice. If no neighbor tries the same candidate color, stick to it, otherwise, just roll the dice again.

This is a simple and efficient algorithm. Unfortunately, its efficiency relies on having fair dices. Scientists call such algorithms *randomized*. Interestingly, many problems (not just the graph coloring problem) can be solved efficiently with randomization. Many of these have in common that the correctness of a solution can be verified efficiently, e.g., you can ask your neighbor for its color to see whether you have the same color. But what if one wants so-called *deterministic* algorithms that do not use dices and are not only efficient if the "gods of the dices" are with us? Over the decades, the quest for efficient deterministic distributed algorithms has become the distributed version of the P vs. NP question: Can every problem whose solution can be verified efficiently by a distributed algorithm also be solved efficiently by a distributed algorithm? The classic P vs. NP question asks the same for non-distributed algorithms. The answer to the distributed question is positive if we allow randomization. In order to also answer it for deterministic algorithms, we have built a large theory that relates randomized and deterministic complexities [100]. Then, building on our theory, researchers from ETH Zürich have resolved the distributed P vs. NP question in the affirmative [167]. In this context, the work of our group laid the foundation for a general derandomization theorem for distributed graph algorithms which is undoubtedly the biggest result in the field of the last decade. Here, our work caused a paradigm shift in how we think about distributed algorithms and distributed complexity theory. Additionally, we have contributed several publications on distributed complexity theory for local constraint satisfaction problems and pioneering works in complexity-theoretic approaches in other decentralized models of computation, e.g., [43, 142, 42].

Communication-efficient distributed algorithms. The primary limitation of network algorithms is that computers within a network can only communicate with their direct neighbors, and there is no central authority overseeing the process. However, in the standard model for such algorithms, i.e., the aforementioned LOCAL model, these messages can be of unbounded size. This is great for developing a rigorous complexity theory but makes the algorithms infeasible in practice. Similarly, the aforementioned "breakthrough results" come with expensive downsides, that is, they rely on a really large communication bandwidth in the network, which is infeasible in practice. Driven by the fact that a few years ago many state-of-the-art algorithms crucially relied on these unbounded message sizes, one research direction that we have pursued in recent years is the study of bandwidth-efficient distributed algorithms, mainly in the so-called CONGEST model [162]. Over the last 5 years, we have contributed to this research direction with 13 publications to date, including the first communication- and time-efficient algorithms for various graph problems. The techniques and viewpoints introduced in these works have been picked up by other researchers and also led to improved algorithms in more general settings without bandwidth restrictions.

Simpler algorithms. State-of-the-art LOCAL algorithms are often very complicated and involved. Besides being communication efficient and fast, one of our objectives is to design simple algorithms. Staying with our graph coloring example, we have shown that small changes to the aforementioned ”roll a dice”-algorithm suffices to exponentially improve upon the previous analysis of the algorithm.

Massively Parallel Computing. In large graph processing frameworks such as Google’s MapReduce or the open-source version of Hadoop a huge graph is arbitrarily distributed on a cluster of machines that communicate with each other in an all-to-all fashion. So, in contrast to the previous setting, there is no relation between the communication network and the input. Surprisingly, insights from network algorithms are very helpful in these frameworks, e.g., our communication efficient algorithms are robust and run in such settings. Recently, our group has also focused on decentralized algorithms in the MPC model and we have contributed several results to build a comprehensive understanding of the field, including pioneering work on complexity theory and state-of-the-art algorithms on benchmark problems such as computing the connected components of an input graph [44].

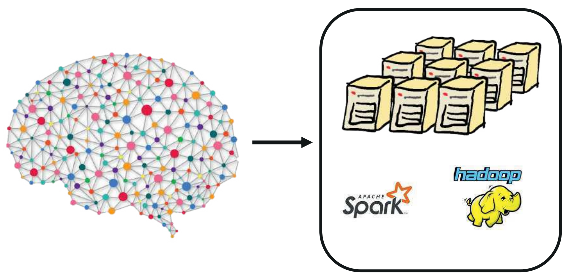

Figure 4.7: The figure shows the massively parallel computing framework to which network algorithms contribute. This figure is licensed under CC-BY.

Several of the group’s works have been awarded substantial recognition by being awarded prizes, best paper awards, or being invited to the special issues of the conference where the results have been published. The latter usually indicates that the work was among the few top papers in the conference of the year.

4.4.3 Group structure and research networks

Currently, the group consists of Ass. Prof Yannic Maus and one PhD student (Manuel Jakob) who is in his first year. Yannic Maus was hired as a tenure-track assistant professor dedicated to one of the five fields of expertise (FoE) of TU Graz. The topic of the FoE is *Information, Communication, and Computing* which perfectly aligns with the research focus of our group. The FoE spans three departments, including the Electrical Engineering, Computer Science, and Math departments. Yannic Maus is also a member of the Graz School of Discrete Mathematics (GSDM).

The group's work is backed up by a strong research network, in particular, with several connections to Israel, Aalto University, and Reykjavik University.

The group is thankful to have received Funding by the FWF. We are planning to hire two new PhD students in 2024 and to significantly extend the size of the group within the next few years. One of these soon-to-be-hired students will also become a member of GSDM.

4.4.4 Outlook

Our future research will build foundations for the emerging algorithmic challenges in our modern world. One long-term objective that we pursue is to build a unified theory and to bridge various computational models in the field. As such we will extend our research to further settings where computation is done with limited information, e.g., we plan the design of streaming and sublinear algorithms.

Most real-world graphs are non-static in nature, e.g., the Facebook graph keeps changing, when new users join or leave the network. We aim to design dynamic algorithms that react quickly to changes without recomputing solutions from scratch. The impact of distributed computing in these areas has been growing and new connections are drawn each day. Historically, the study of dynamic algorithms has been strong, particularly in Austria, and extending our research to that field will improve the connections between TU Graz and other research institutions in Austria.

TU Graz itself offers ample possibilities for further connections with other groups. For example, we recently started a collaboration with the group of Michael Kerber from the Institute of Geometry in the Math department. He is a specialist on topology, and we will use topological methods to rule out the existence of certain distributed algorithms.

In general, distributed computation has further additional connections to almost all aspects of computer science, e.g., communication cost is a driving cost factor in training neural networks, large dynamically evolving networks evolve in the area of communicating autonomous vehicles, etc. Connections reach as far as biomedical engineering; there are several strong groups in our department whose main focus is to understand the brain. In all of these settings, many entities work together to reach some big goal and there are natural connections to distributed computing. For the next few years, the group will focus on its timely core objectives, but we are always ready to adapt to our modern world.

4.5 Group Schweiger - Intelligent Energy Systems

The interdisciplinary research group Intelligent Energy Systems was founded in 2018 and conducts research in computational methods and semantic data related to intelligent buildings and energy systems. Together with the research group of Katja Corcoran from the University of Graz (Institute of Social Psychology), the group also conducts research in the field of active user participation.

The motivation for the research arises from social and political premises: We need to reduce energy usage and increase the share of renewable energies. As buildings are responsible for 40 % of final energy consumption and 30 % of total energy-related CO2 emissions, they are one of the critical factors and levers that need to be taken into account. In recent decades, traditional engineering disciplines have optimized systems and components for buildings and energy systems. Nowadays, we have the most significant potential to reduce energy usage by establishing intelligent, flexible systems that orchestrate the interaction between energy supply from volatile sources such as wind and solar, various storage technologies, and the energy demand side (e.g., buildings, industry). Buildings have a central role: they must transform from static to active, intelligent entities whose energy usage must be significantly reduced and whose demand can be flexibly adapted to external conditions such as prices and energy availability. Studies show that the energy usage of office buildings can be reduced by 30 % [66] through predictive control and by up to 15 % through the introduction of automatic fault detection [61].

Many steps towards intelligent systems have one thing in common: problems arise in various disciplines and industries. Many solutions require interoperable systems and scalable, robust services. To avoid getting stuck in buzzwords and the inflationary use of terms like ”intelligent building,” ”smart city,” and ”smart grid,” computer science needs to be intensively involved in shaping these research areas.

4.5.1 Group Structure

The group has raised more than 4 million euros in funding in the last five years. This includes applied projects in the field of intelligent buildings (e.g. H2020 Smart2B, EraNet I-Greta), projects in the field of computational methods (e.g. FFG DomLearn, Zukunftsfond DISTEL), and projects in the field of active user participation (e.g. FFG GameOpSys, Zukunftsfond ANSERS). The interdisciplinary character of the group is also reflected in the fact that the group carries out projects with institutes from 5 faculties of Graz University of Technology.

4.5.2 Scientific Highlights

One of the group's research areas is Co-Simulation. Co-simulation is a generalized form of simulation where a coupled system is simulated through the composition of different simulation units. One needs a co-simulation scenario and an orchestrator algorithm to run a co-simulation. Each simulation unit is seen as a black box, capable of producing outputs and consuming inputs, according to the model it represents. In the NextHyb2 project, methods for coupling machine learning models with classical physics-based models were developed [200, 76]; in the DigitalEnergyTwin project, co-simulation tools for applications in printed circuit board manufacturing were developed [170, 138]. The group has also published empirical work related to Co-Simulation standards, tools, and challenges [173, 37].

Together with colleagues from the Institute of Thermal Engineering (Thomas Mach) and the Building and Technology Department of Graz University of Technology (Siegfried Pabst), the group has launched the research initiative ”Innovation District Inffeldgasse”. This initiative is an integral part of the climate-neutral TU Graz 2030 plan and is intended to bundle research activities in the field of intelligent buildings and energy systems. Within the Inframonitor project, the group developed the IoT and data infrastructure for the Innovation District [168, 38]. Based on this data, the group has developed various data-driven energy services [199, 50, 171, 169].

4.5.3 Impact Highlights

In 2022, Gerald Schweiger and Franz Wotawa founded the spin-off DiLT Analytics. The spin-off aims to convert the developments of the two research groups, Intelligent Energy Systems (Group Schweiger) and Software Technology and Artificial Intelligence (Group Wotawa), into solutions for the intelligent energy system of the future. The founding team was recently expanded to include the two (former) IST employees Thomas Hirsch and Thomas Schwengler. DiLT currently has eight employees from the fields of computer science, physics, building technology, and control engineering.

4.6 Group Slany - Agile Software Development, Mobile Applications, and Didactics of Programming

4.6.1 Research Focus

The group has a rather broad range of research foci, which also changed over time, as reflected in the broad range of publications, funded projects, and interests of the group's members. While the focus initially was on theoretical and practical combinatorial optimization problems and algorithms, including artificial intelligence algorithms such as genetic algorithms and expert systems, the group's interests quickly turned to the, in the early 2000s, emerging agile software project management methods, in particular test driven and behavior driven development methods, applied to the then also increasingly popular mobile platforms — initially for Nokia's Symbian platform, later for Android, Windows Mobile, and iOS. Starting in 2009, the group also turned towards the teaching of computational thinking methods, in particular programming, and as there was no mobile platform for doing so, but young people increasingly owned smartphones, all these aspects were combined in the Catrobat project initiated in 2010, which has been, in all its aspects, the main focus of the group, ranging from didactics, psychology, sociology, gender studies, security, DevOps, open source philosophy, up to programming language design, hardware testing, Internet of Things, and user experience. According to OpenHub, more than 800 person-years-of-effort have been invested in the project, with nearly 3 million lines of code and 10 million users worldwide. In recent years, the focus has returned to AI-based systems, in particular their integration into mobile applications for coding and into their support of the learning of coding. A completely new research interest since 2020 has been the interdisciplinary investigation of the cognitive capabilities of cephalopods and their similarities to how Large Language Models represent knowledge and learn new skills, which is being conducted collaboratively with marine biologists.

4.6.2 Group Structure

Besides Slany, currently, two university assistants are working full time in the group, Stefan Kutschera (security and privacy aspects) and Patrick Ratschiller (DevOps), with Sarina Gursch working as a research assistant on Gender studies and having just completed her PhD in the framework of several FFG funded projects and a Horizon2020 project.

4.6.3 Scientific Highlights

Our recent paper "Perspectives on gender mainstreaming in international co-operation in STI: A comparative study", International Conference on Gender Research 6 (1), 146-154, received the Mind the Gap Diversity Award 2023 from Graz University of Technology. Another maybe noteworthy paper, among a plethora of others, was our 2018 paper on "Analyzing and managing complex software ecosystems: A framework to understand value in information systems" published in IEEE Software 36 (3), 55-60.

4.6.4 Impact Highlights

The Catrobat project has had a large impact on coding skills for nearly 10 million youth in the age range 7 to 18 around the world, with some of them spending a decade using our apps to code their own apps on their smartphones. This has been reflected in nearly 20 prizes, awards, and other recognitions: The project was selected 10 times as a Google Summer of Code mentor organization, won the Austrian National Innovation Award for Multimedia and e-Business in 2013, the Internet for Refugees Award 2016 issued by Internet Foundation Austria, the Re-Imagine Education Award Europe 2016 issued by Wharton School of the University of Pennsylvania, up to the Huawei Best Social Impact App Award in 2022.

4.7 Group Steinbauer-Wagner - Autonomous Intelligent Systems

4.7.1 Research Focus

The Autonomous Intelligent Systems (AIS) group is interested in the methods, tools, and architectures that are needed to build autonomous intelligent systems such as agents and robots. In particular, we are interested in the complete decision-making process (perception-cognition-acting) for such systems and in the related challenges that arise if such systems are deployed in uncertain real-world environments. Due to the non-determinism in the systems, the environment, and their interaction advanced methods for making and executing decisions are needed to obtain reliable systems we can trust in. In order to achieve this goal AIS resides in the intersection of Software Engineering, Artificial Intelligence, and Robotics.

Research-wise the AIS group is active in comprises architectures for making and executing decisions, action planning in various forms, supervision and diagnosis of system components and behaviors, robot navigation in indoor and outdoor environments, and explainable and trustworthy human-robot interaction. AIS is working on integrated robot systems for various application scenarios such as production, logistics, planetary exploration, and disaster response.

The core research topic is **Dependable Autonomous Systems**. Autonomous systems, such as mobile robots, must integrate decisional autonomy and dependability to function effectively and reliably in complex and uncertain environments. Unlike structured domains like aviation or automotive engineering, these systems operate in open-ended settings where tasks and environmental conditions are often partially known or unpredictable. This introduces uncertainty and non-determinism at multiple levels, including sensing (due to sensor noise and limitations in scene understanding), reasoning (due to incomplete or ambiguous knowledge), and action execution (due to unforeseen interactions with a dynamic environment). As a result, ensuring reliable performance while allowing flexible, autonomous decision-making remains a significant challenge.

A major limitation of current autonomous systems is their incomplete adaptability and knowledge representation. While machine learning has significantly enhanced perception and control, its narrow scope and lack of explainability make seamless integration with high-level reasoning difficult. Robots still struggle with unexpected environmental changes, open-ended tasks, and a lack of common sense reasoning, often leading to performance degradation or failure in novel scenarios. To improve reliability, model-based approaches such as diagnosis (DX) and fault detection and isolation (FDI) have been explored, addressing failures in hardware and software as well as in the perception-decision-action cycle. However, these methods are still constrained by the complexity and unpredictability of real-world interactions, requiring more advanced techniques for real-time supervision, adaptation, and self-repair.

Achieving long-term autonomy — where systems operate for months or years without human intervention poses additional challenges. Over time, both the robot (e.g., component wear, sensor degradation, software drift) and the environment (e.g., seasonal changes, new obstacles, shifts in human activity) evolve, requiring continuous adaptation. The system must detect deviations, determine whether they represent transient noise or significant changes, and update its models accordingly. This requires advanced reasoning capabilities, combining semantic understanding, hybrid knowledge representation (qualitative and quantitative), and efficient planning under uncertainty. Furthermore, autonomous systems must be able to handle completely novel situations with only limited real-time data, which remains an open challenge for both logical inference and machine learning methods.

Future research must focus on developing hybrid reasoning approaches that integrate multiple knowledge representations and learning methods while ensuring real-time efficiency. Additionally, improving the resilience and robustness of autonomous systems requires more advanced techniques for self-monitoring, adaptive planning, and fault-tolerant operation. Addressing these fundamental challenges will be crucial to enabling dependable, intelligent autonomous systems capable of long-term operation in complex, open, and uncertain environments.

Long-term Autonomy and Model Adaptation Humans are very good at anticipating changes in a task or environment as well as the need for new or adapted skills and in learning new skills with a few examples and interactions. This capability is essential for autonomous systems that serve in open tasks and environments for a long time without human supervision [154, 152]. It appears quite unrealistic that for a complex system, task, and environment every relevant aspect can be modeled or learned in advance. Similar to children the system will start with incomplete models or skills and need to extend or refine it online. Here we need to distinguish the different ingredients of the models. First, given that there is a vocabulary (e.g. symbols, predicates) that can describe all relevant concepts, skill models (e.g. action descriptions using preconditions and effects) can be rather simply adapted for missing effects and preconditions [174, 69]. There are works using simulation intensively, but learning such information from a few interactions is still an open challenge. If the basic vocabulary is still sufficient, the integration of novel unseen objects into a domain description is moderately tricky from the theoretical perspective. One will need concepts in the knowledge base to describe such objects and meta-actions to generate them in the knowledge base. The most interesting case is when concepts to describe relevant phenomena and symbols to represent new skills and objects are missing. In the first place, the system needs to have the capability to realize that lack which means making a distinction between a goal failing because of some internal or external reason, because a skill is missing, or because some concept is not represented. Here interaction with reinforcement learning is needed as the low-level skill needs to be learned as well as with techniques from Deep Learning like autoencoders which can help to bridge the symbol grounding gap.

Temporal Planning and Plan Execution Automated planning has been of interest in AI research for decades. A huge corpus of work was conducted for the modeling of planning domains and for obtaining plans for given goals efficiently. This declarative approach to problem-solving is among the computationally hardest problems in Artificial Intelligence and can reflect various properties such as discrete and continuous change, non-deterministic effects, or temporal aspects. For a long time, researchers in planning were focused on plan generation mainly. Recently, there has been a growing interest in integrating planning into intelligent systems and focusing on plan execution. A prominent example is the ROSPlan system. Executing plans on a robot in a complex environment posts interesting research questions, as the planning process uses assumptions about the system and environment, domains are non-deterministic, and the fidelity level of domain descriptions is limited. To be able to use temporal planning in flexible production settings where temporal relations, duration, and deadlines play a role we developed a framework for temporal planning in plan execution [53, 55]. Key contributions are the generation and use of macro-actions in temporal planning and automated adaptive goal reasoning to deal with the complexity of temporal planning and an execution framework that can detect logical and temporal discrepancies in plan execution early [52, 54]. So far this detection often leads to simple re-planning losing valuable information from the previous plans. Future work is needed in the direction of effective repair of temporal plans. Currently, most well-performing temporal planners use state-space search. In contrast to plan-space search, the underlying representation is less suitable for efficient plan repair. Thus, planners based on plan-space search like FAPE and their integration into temporal plan execution need to be investigated.

Hierarchical Supervision In general, autonomous systems use layered architectures for decision-making involving an abstract deliberation layer (e.g. planning), an executive layer for action refinement (e.g. behavior trees), and a skill layer (e.g. controllers and policies). Although we developed already various monitoring and diagnosis methods for the individual levels [153, 151, 163], there is an interdependency between these layers when it comes to execution monitoring. There are situations that cannot be detected with information of one layer, rather need to integrate information of different levels in monitoring and reasoning. Due to the idea of abstraction and refinement, the temporal context is enlarged when moving up the hierarchy while the depth of information is reduced and vice versa. For instance, on the level of skills, there is rich information on the state of the skill but the temporal context on possibly relevant previous executed skill is missing for a decent fault diagnosis. On the contrary, the sequence of skills is available for monitoring on the higher levels, but there might be relevant details missing. This dilemma is often solved ad hoc using specific domain knowledge and hand-crafted rules. To utilize the hierarchical structure better for supervision in the future a general theory of hierarchical supervision needs to be developed. This development will be based on existing hierarchical diagnosis approaches for hardware, hierarchical action models, and the theory of abstraction/refinement and its operationalization.

Automated Navigation in the Wild During the last years, we developed an integrated approach for automated navigation for remote off-road environments like alpine areas. Possible application areas include inspection of alpine protection structures, support of mountain rescuers, forestry, agriculture, mining, and remote logistics. Key challenges in navigating such environments are that in contrast to on-road automated driving detailed maps are missing, there is an insufficient environment understanding of natural scenes, and the terrain and topology posts complex control problems for the navigating vehicle. To address the first challenge, we developed an automated pipeline for the generation of costmaps from earth observation data. Using machine learning and information like orthophotos, digital height maps, and land cover a traversability analysis is done. Using data from UVAs the spatial and temporal resolution of the costmaps can significantly be improved [70, 73]. By using short defined locomotion recordings on various terrains and slopes, robot-dependent costmaps can easily be obtained [74]. Environment perception is mainly based on geometric reasoning on 3D point clouds and the segmentation of camera images. The navigation architecture uses a stack of a search-based global planner, an optimization-based local planner, and a controller for trajectory-following. The novelty here is that instead of cyclic planning, trajectories are only generated when a new goal arrives or significant changes in the environment are detected [114]. This ensures smooth and efficient navigation. Open issues for the future are direction-dependent costmaps (up-hill versus down-hill), optimization-based planners that consider the 2,5D topology and complex vehicle-terrain interaction, and an improved perception using multi-modal sensors that can cope with complex vegetation and diverse weather conditions.

Transparent Shared-Autonomy Although we discussed here so far autonomous systems there are reasons for and room for semi-autonomous systems. There will be always situations where the autonomous system will reach its limits in reasoning, sensing, and action. Here in the sense of leveled co-operation between the human and the machine, human ingenuity is used to complement the limited reasoning, sensing, and acting of the machine (e.g., template-based selection of an object for manipulation, manual steering through a narrow passage, provision of a better path) [98]. To improve trust in the system by the human and to reduce cognitive load in using it, transparency is of immense importance. The system needs on the one hand to be transparent about its activities, but should also be able to give a transparent explanation when delegating activities back to the user or when asking for support. To do so, the system needs to be able to reason about its state and capabilities and to communicate shortcomings understandably [71]. In the domain of assistance robots for first responders, we developed a method based on consistency-based diagnosis to detect and communicate failures in motion planning [72]. Experiments with human subjects confirmed that this approach improved trust and lowered cognitive load. In the future, a more elaborated method for the estimation of the actual confidence level of the autonomous system about its current capabilities needs to be developed. While this was rather easy to achieve for constraint-based path planning, it remains unclear what representation and reasoning are suitable for complex activities executed by the robot.

Educational Robotics and AI Literacy Last but not least the research addresses the outreach to the public and the improvement of STEAM and in particular Computer Science education in schools because modern motivating and appropriate teaching of CS, Computational Thinking, and recently AI is the basis of a broad understanding of the technology we develop and the recruiting of motivated skilled students in the future.

In the last almost 2 decades the group developed many didactical concepts and teaching material for Educational Robotics where the robot is used as a tool to teach CS as well as robotics concepts [68, 123]. Moreover, the group has been a driving force in the establishment of the RoboCupJunior initiative in Austrian schools and the preparation of ongoing or in-service teachers to use the concepts in classes. An important aspect has been from the beginning the sound scientific evaluation of the activities, which was always a weak spot in this domain [124, 122]. Recently the focus shifted towards the topic of AI Literacy in schools [123]. In the development and evaluation of concepts and material for teaching AI to young learners the group has a holistic view of AI and a strong focus to also include classical symbolic methods in teaching as well as on understanding the underlying concepts rather than only use AI tools [176].

4.7.2 Group Structure

The research group for Autonomous Intelligent Systems (AIS) is headed by an Associated Professor and comprises a varying number of PhD students funded by projects of national funding agencies (FWF, FFG) and various strategic projects such as the Disaster Competence Center Austria (DCNA), the PhD college for Innovation in the Austrian Education Systems, or the Lead Project on Dependable Systems. The research group also hosts a high number of master students working on their thesis in funded or strategic internal projects. The group also hosts a full robotics lab with several robot platforms and an advanced software framework which is the basis for high-class teaching in Robotics and Intelligent Systems and the conducted research projects. Application domains range from classical service robots and logistics over robots for disaster response and agriculture to deployments of robots in unstructured environments such as alpine regions and forests.

The AIS robotics lab also hosts two successful student teams of the Graz University of Technology. The team TEDSUAR (TU Graz Field Robotics Team) has been active since 2009 and was part of several competitions (Winner RoboCup 2016 Autonomy RoboCup Rescue Robots and Elrob 2022 Mule) and field deployments (Mars Analog Missions in Oman, Israel, and Armenia or Austrian Alpine Robot Trials). The team GRIPS (Graz Roboust Intelligent Problem Solver) masters the application of Cognitive Robotics to the Flexible Production Challenge and is a 3-times world champion in the RoboCup Logistics League.

AIS is also the base of the RoboCupJunior activities in Austria and organizes the annual RoboCupJunior Austrian Open tournament and teaching activities for schools teacher in the domain of Educational Robotics and AI.

4.7.3 Scientific Highlights

One of the scientific highlights of AIS is the successful integration of advanced monitoring, diagnosis, and repair capabilities into real robot systems to significantly improve the dependability of autonomous intelligent systems. This includes on one hand the development of methods for the different aspects of an autonomous system (hardware, software, perception, cognition, and behaviors) [177, 179] as well as the integration of them into a full robotics software architecture [180, 210, 151].

Another scientific highlight of AIS is the development of a full-fledged navigation architecture for unstructured complex off-road environments such as alpine regions or forests. The developments start from considerations for a flexible suitable navigation structure [114] and range over the development of traversability maps [70, 73, 74] and the environment understanding [65] to trajectory planning and execution [75]. This navigation framework has been the basis of several research projects in areas such as disaster response [99, 64], planetary exploration [121], or military logistics [41].

4.7.4 Impact Highlights

Two successful startups were founded by alumni of the AIS research group. The company incubedIT was funded by the former 8 members of the TU Graz RoboCup Simulation Soccer 2D team. The company develops software solutions for robot navigation and fleet management in warehouse automation. The company was recently acquired by Verizon. The company ARTI – Autonomous Robot Technology was founded by former members of the TEDUSAR RoboCup Rescue Robot team. The company develops software solutions for the automation of robots in various domains such as logistics, inspection, or large outdoor logo printing. There is ongoing cooperation with the companies by recruiting graduates from the lab, joint supervision of master's theses, or funded research projects. These facts highlight the high standard of teaching in the group and knowledge transfer into the industry.

In 2009 Franz Wotawa and Gerald Steinbauer-Wagner co-organized the RoboCup 2009 in Graz. RoboCup is one of the largest and long-lasting initiatives to promote education, research, and development in Robotics and AI. Making Graz the capital of Robotics for a week fostered the visibility of AI-driven Robotics in Austria and set the nucleus for a vibrant Robotics network in Austria and several follow up activities.

4.8 Group Vogtenhuber - Discrete Mathematics & Theoretical Computer Science (DMTCS)

The group performs foundational research in the areas of discrete mathematics and theoretical computer science, with a focus on structural, algorithmic, and complexity-theoretic questions in discrete and computational geometry as well as graph drawing and network visualization. As such, it bridges mathematical and computational viewpoints on discrete structures.

In addition to research, the group is strongly involved in teaching from the undergraduate to the PhD level. This includes regularly performing courses on algorithms and data structures, the design and analysis of algorithms, discrete and computational geometry, complexity theory, and several seminars, as well as the supervision of Bachelor-, Master-, and PhD-theses.

The group was created in 2013 at Graz University of Technology via a new tenure track position between mathematics and computer science, funded by the Faculty of Mathematics, Physics, and Geodesy as well as the Faculty of Computer Science and Biomedical Engineering.

4.8.1 Research Focus

Two fundamental structures at the core of discrete mathematics and theoretical computer science are arrangements of geometric objects and (drawings of) graphs. They are of great importance in both computer science and mathematics, and they serve as crucial tools in applications. Many problems of geometric nature deal with arrangements or graphs. This includes several prominent fundamental open questions in discrete mathematics (such as the number of triangulations or empty triangles spanned by a set of n points, or the existence of plane substructures in drawings of graphs), as well as the search for efficient algorithms for geometric problems (such as finding shortest routes in a road network or a public transport network). Moreover, many seemingly unrelated questions can be formulated in terms of arrangements or graphs. Hence the successful understanding of these structures can impact a wide variety of problem domains.

Our main focus is structural, algorithmic, and complexity-theoretic questions on graphs and their drawings as well as arrangements of geometric objects. Particularly, we work on combinatorial properties of point sets, graphs, and different types of graph representations, abstractions of graph drawings (such as order types and rotation systems), and algorithmic and complexity-theoretic questions on these structures. We give some examples in the next section.

4.8.2 Group Structure

The group consists of Assoc.Prof. Birgit Vogtenhuber and a varying number of third-party funded PhD students as well as associated Master and Bachelor students. As usual for associate professorships at TU Graz, the group has no further university-funded positions. The group is thankful to have received and received funding for positions by the FWF.

Birgit Vogtenhuber holds a habilitation in theoretical computer science. She has up to now supervised and co-supervised 5 PhD students (one of them ongoing). Prior to founding this group, she obtained a diploma in technical mathematics and a PhD in computer science. In addition to her natural involvement in the Doctoral School of Computer Science at TU Graz, she is a member of the Doctoral School Mathematics and Scientific Computing and a founding member of the doctoral program *Graz School of Discrete Mathematics (GSDM)*. Further, she is part of the TU Graz field of expertise *Information, Communication, and Computing*, which focuses on the challenges of the information age and fosters interdisciplinary connections between electrical engineering, computer science, and mathematics.

The group's work is supported by a strong international research network with connections to researchers in Germany, Switzerland, the Czech Republic, Spain, Mexico, and several other countries. In particular, Birgit Vogtenhuber has been PI of joint research projects with ETH Zürich, TU Berlin, and FU Berlin, as well as Charles University, funded by the FWF and the OEAD, respectively. Within TU Graz, the group collaborates with researchers from computer science (including the groups of Oswin Aichholzer and Yannic Maus) and from mathematics (especially within the GDSM). We are happy to announce a recently granted highly competitive FWF doc-funds project within the GDSM that enables funding for 13 PhD students (one of whom will join this group).

4.9 Group Wallner - Knowledge Representation

The group specializing in computational aspects of knowledge representation and reasoning (KR&R) was recently established at Graz University of Technology. With roots in both the Technische Universität Wien (TU Wien) and the University of Helsinki, this group started its work in Vienna at the Institute of Logic and Computation. In 2020 the group leader started a tenure-track position at the Institute of Software Technology.

The field of KR&R is part of the broader field of Artificial Intelligence (AI) and is featured in many conferences and journals in Computer Science. Broadly construed, KR&R is concerned with foundational research on how to represent knowledge and how to perform reasoning, based on knowledge. Tightly connected with both representational issues and reasoning aspects is formal logic, whose aim is to formalize rational reasoning. Logical reasoning exhibits high computational complexity, witnessed, e.g., by the first problem to be shown to be NP-complete: the Boolean satisfiability problem (SAT), which is at the heart of many processes of logical reasoning and declarative problem-solving.

4.9.1 Research Focus

The research and teaching of the group focus on

- formal studies of prominent logic-based representations of knowledge, and
- addressing challenging computational reasoning tasks arising in KR&R.

In brief, our research agenda is to further understanding of complex forms of reasoning in knowledge representation, and to bring promising approaches closer to application, by going from theory to practice. In our research, we, in particular, focus on complex reasoning problems arising in KR&R, which are naturally to be found in KR&R due to connections of the field to formal logic.

The main methods we employ can be summarized as follows (see also Figure 4.8).

1. Providing foundational insights into problems targeted: establishing fundamental theoretical properties, e.g., performing complexity analysis.
2. Drawing from and extending the state of the art of algorithmic approaches and systems for hard problems: for instance utilizing algorithms and systems for Boolean satisfiability and answer set programming [58].
3. Performing experimental evaluations: prototype implementation, feasibility studies, and experiment design.

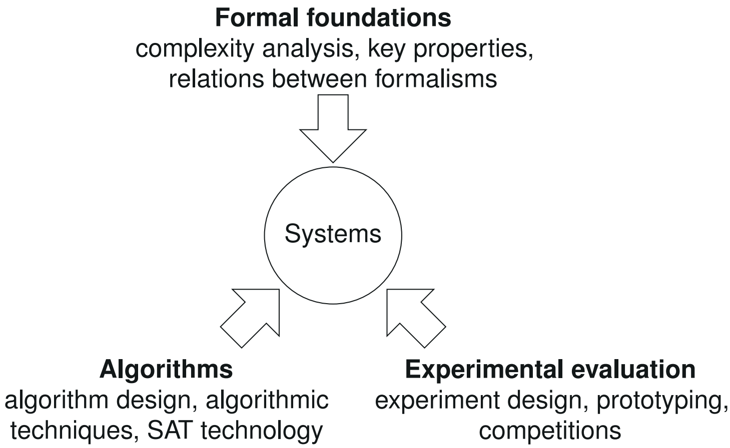

Figure 4.8: Methodological approach to complex reasoning tasks.

That is, towards reaching the overarching goal of development of systems for important problems in AI & KR&R, a thorough understanding of *what* is to be solved is required. During the course of our research, we provided fundamental results for several heterogeneous and complex problems in AI, in particular focusing on computational complexity analysis, which gives key insights into the problems at hand.

Inherent to problems with roots in logic comes a computational price: many problems within AI and KR&R are NP-hard, or even “beyond” NP (in the formal sense of being hard for a class assumed beyond NP, e.g., within a higher level of the polynomial hierarchy). Intuitively, this means that naive algorithms are unlikely to lead to efficient systems. This calls for sophisticated algorithmic design. Due to the advent of advanced systems for NP-hard problems, such as solvers for the Boolean satisfiability (SAT) problem, modern solvers provide an efficient means of tackling certain archetypical problems. In our research, we provided and contributed to several algorithmic solutions for problems in KR&R situated at different complexities via the usage of modern (SAT) solvers. Based on our experience, it is our firm belief that a deep understanding of problem complexity, algorithms, and available systems (such as SAT solvers) are prerequisites for efficient system approaches to hard problems.

Complementing theoretical underpinnings, and experimental evaluations to show the feasibility of implementation approaches are indispensable.

The three key approaches are strongly interconnected to reach our vision of systems solving intrinsically hard problems in AI. For instance, complexity analyses lead to algorithmic design and, based on that, also experimental design; however, studying algorithmic approaches and implementations can lead to new insights and novel research questions that, in turn, have to be investigated on a fundamental basis.

4.9.2 Group Structure

The group began its work in Vienna, at TU Wien, in 2017. During the time in Vienna, Adrian Haret and Zeynep G. Saribatur, alongside the group leader, worked on a project funded by the Austrian Science Fund (FWF). Haret finished his PhD studies in 2020 and by now is an assistant professor at the Munich Center for Mathematical Philosophy (MCMP). Saribatur is, at the time of writing this chapter, a Postdoc at TU Wien and holds a Hertha-Firnberg grant from FWF.

After moving to TU Graz at the end of 2020, the group recruited two PhD students with a new FWF grant: Iosif Apostolakis and Andrei Popescu. Both PhD students are currently in their first year.

The group is thankful to be able to carry out their foundational research with funds from the FWF.

The webpage of the group can be found at

`www.tugraz.at/institute/ist/research/group-wallner`.

4.9.3 Scientific Highlights

This group focuses its research on complex reasoning tasks that can be found in the broader field of KR&R. Summarizing, we contributed to diverse fields in, or related to, KR&R: formal argumentation [46], belief change [36], approaches to handle and measure inconsistency [102], and computational social choice [57], utilizing state-of-the-art approaches to solve computationally complex problems such as Boolean Satisfiability and answer set programming. In this section, we overview some research of the group in formal argumentation.

Since its inception a few decades ago, formal argumentation has established itself, within the broader field of AI, as a versatile form of non-monotonic reasoning that provides foundational approaches for arguing in favor, or against, claims that are under scrutiny [49]. Approaches to argumentation find applications in a heterogeneous landscape [39], e.g., argumentation finds applicability for several domains connected to debates, justification, or resolving conflicting opinions, such as providing justifications of claims in legal reasoning, means to compare treatments in medical sciences, or ways to engage in dialogues for proposed policies in e-government tools. Generally speaking, within the broader AI research, argumentative reasoning presents itself also as an approach to *explainable AI* (XAI)—a current hot and largely unresolved topic within AI [62].

With an inherent goal of automated reasoning, computational models in formal argumentation provide reasons for supporting a certain claim based on a construction of (counter-)arguments in favor, or against, the claim in question [46]. The starting point is a knowledge base comprising diverging opinions and conflicting information, e.g., in a rule-based form, from which

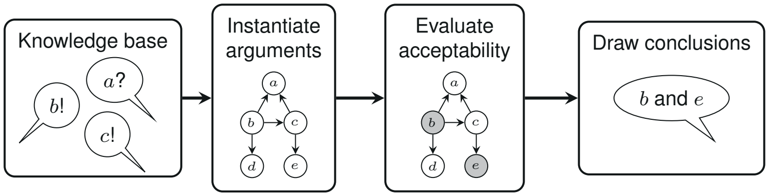

Figure 4.9: The argumentation process [60].

arguments are drawn. Support of a queried claim can be shown in various ways depending on the intended application. A prominent approach is based on forming admissible sets of arguments. Under this view, a queried claim is warranted if an argument concludes that the claim can be defended against all counterarguments, possibly utilizing further defending arguments that are non-conflicting. In this way, a dialectical explanation and justification for supporting the claim is provided, using formal argumentative machinery to analyze reasons against and reasons in favor. This kind of "workflow" is illustrated in Figure 4.9.

Formalisms in the field of argumentation provide principled approaches to represent arguments and reason on argument structures. *Structured argumentation* refers to a collection of approaches to construct arguments and reason on arguments, while *abstract argumentation* refers to the case when reasoning is performed on an abstract representation of arguments, typically consisting of arguments as abstract entities, with only relationships given between arguments. Operating on abstract arguments and their relations suffices to reason in an argumentative way for a variety of application scenarios.

Our research contributed significantly to the overall area of computational argumentation. In more concrete terms, we advanced the state of the art in the following three directions.

1. Reasoning in advanced argumentation formalisms.

2. Dynamics of argumentative reasoning.

3. Supporting argumentative explainability.

With dynamics, we mean here dynamically changing states of argumentation, such as during a debate.

In all three directions, we laid out or expanded the fundamental complexity landscape of many core problems within argumentation. There are various ways of representing arguments—reflecting on the heterogeneity of argumentation in general. We contributed to several such formal representations, such as argumentation frameworks (AFs) [67], abstract dialectical frameworks [59], assumption-based argumentation [51], and ASPIC$^+$ [146].

Let us give an overview of two of our works. For argumentative and dynamical reasoning, we looked at the so-called enforcement problem [47] and its computational properties [193]. Briefly put, in the enforcement problem, within formal argumentation, one aims to find ways of "arguing in one's favor". For instance, in the simple yet effective way of representing arguments as abstract vertices and considering a directed "counter-argument" relation, one can look for additions, or general modifications, to the underlying graph structure to achieve a certain argumentative goal.

Consider a simple scenario shown in Figure 4.10. In the left-most directed graph (argumentation framework F_1), we see a representation of five abstracted arguments a to e, which can stand for arguments that counter another according to the directed edge relation. For instance, argument a "attacks" argument b. One can interpret such directed argument graphs, for instance, as follows. Since there is no reason to the contrary one can accept argument a (and its contents). This is because no argument stands against a (no counter-argument to a exists in F_1). Since we can deem to accept a, we should not accept b, because we accept a direct counter-argument to b. In turn, we rejected the counter-argument of both c and d (that is, b), and have no reason (anymore) to reject them. Likewise, e is deemed rejected, since both counter-arguments of e (c and d) are deemed accepted.

Say, we desire to argue in favor of argument e. In the current state of the debate (we are assuming F_1), there is no hope for accepting e. However, we can find ways of posing new arguments to attack just the right arguments, which can turn e being acceptable. For instance, posing arguments x and y to attack c and d would do the trick: then all counters to e are argued away (F_2). Similarly, one could look for ways of arguing against a, by considering the scenario in F_3: expanding with argument z to "defeat" a (and turning the whole acceptance chain from above the other way around).

The enforcement problem can be defined as a computational problem of finding ways of arguing in favor of a certain goal (such as finding a specific argument acceptable). Since the formalisms to represent argumentation come in many flavors, so does the enforcement problem, we surveyed several of them in the recent second volume of the Handbook of Formal Argumentation [48].

In our work, we look at the complexity of enforcing acceptance of arguments, under several parameters and scenarios [193]. It turns out that many of these are intrinsically computationally complex, i.e., are NP-hard. To overcome this complexity barrier, we looked at ways of developing prototype implementations for these problems utilizing Maximum Satisfiability Solving (MaxSAT) [40], which in our evaluations showed promising run-time performance.

In a different direction, we also looked at a more advanced formalism to represent argumentation scenarios, in particular one that expands on the relations between arguments and allows more flexibility in specifying the relations. In abstract dialectical frameworks (ADFs), one can specify a Boolean formula per argument to state acceptability. For instance, in Figure 4.11, there are four (still abstract) arguments. Argument c has an associated "acceptance condition" of the form $\neg a \vee \neg b$, specified in Boolean logic, which intuitively represents the condition that c is deemed acceptable if a is not accepted ($\neg a$) or b is not accepted ($\neg b$).

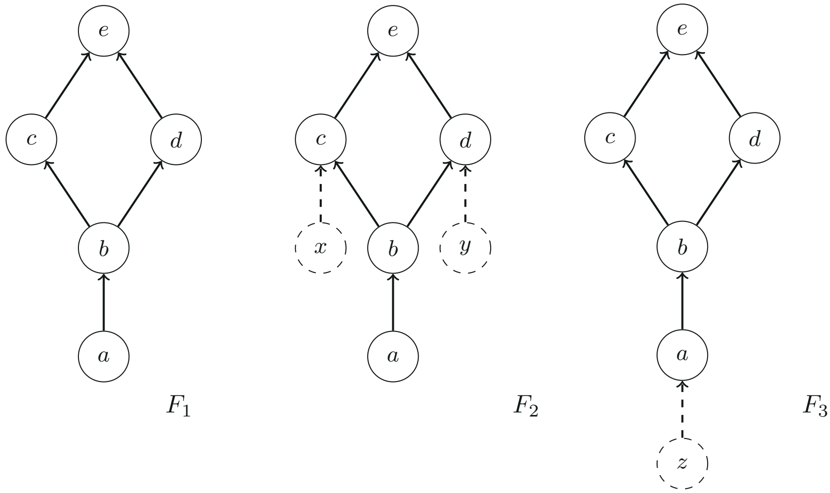

Figure 4.10: “Enforcing” to accept argument e.

Inspired by the use of ADFs for legal reasoning [35], we looked at the computational complexity and algorithm design for reasoning in ADFs, this time in a static way, i.e., when ADFs do not change but are static. We showed earlier [183] that even static forms of reasoning lead to complex problems in ADFs. In a more recent work [141], we looked at so-called “fragments” of the general ADF language. Or, in more direct terms, when considering restrictions on ADFs, we aim to answer the associated question of whether reasoning becomes computationally simpler.

We looked at several restrictions to ADFs, e.g., properties of the underlying graph structure such as acyclicity, bipartiteness, and symmetry. For instance, acyclicity might occur if arguments are posed after each other and previous ones do not counter newer ones (which, however, is not satisfied in all cases). Bipartiteness can occur in a two-party debate, where arguments of each party only relate to arguments of the other party. In addition to these, we looked at restrictions that relations between arguments have a somewhat simpler form and further restrictions that, in a sense, make an ADF more “deterministic”. It turns out that only rather strong assumptions, like acyclicity, lead to mild complexity, in the form of polynomial-time algorithms, for many tasks, whereas many of the other forms, although simpler in cases, still exhibit NP-hardness. We utilized these insights gained from complexity analysis and designed algorithms based on these results that work well if the given ADF is in one of the fragments, or is close to one of the fragments. Our experimental evaluation showed that this approach outperformed, at the time of publication, the other competitor systems for ADF reasoning.

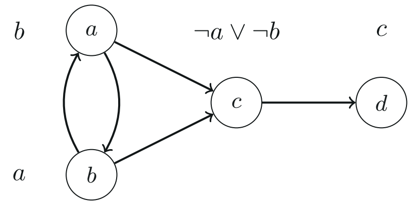

Figure 4.11: An abstract dialectical framework.

4.9.4 Outlook

Given a brief summary of the research of our group, we think that rigorous formal foundations for complex reasoning tasks, such as those that arise in KR&R, are indispensable to the development of efficient systems for such reasoning. While complexity analysis itself can only be one step in this direction, we think that this step can significantly expand insights and further algorithm design. Utilization of adequate (in complexity terms) systems such as systems for the Boolean Satisfiability problem can be a benefit, since implementation does not need to start from scratch, but can use decades of research and engineering that went into such solvers.

In formal argumentation, our works can be the basis, for instance, for pushing the formalizations of argumentation towards practice. For instance, ADFs were employed for legal reasoning [35] and more recently we showed foundational insights into related formalisms currently prototyped at the Dutch National Police [157], which can also lead to extend the reach of such prototypes.

4.10 Group Wotawa - Software Engineering and Artificial Intelligence (SEAI)

4.10.1 Research Focus

The software engineering and artificial intelligence group (SEAI) focuses on applying software engineering to artificial intelligence applications and using artificial intelligence for software engineering challenges. The latter include the automation of software testing and automated debugging. For both research directions, members of SEAI have provided discussions and solutions that have been published. Of course, without collaboration within the group but also with colleagues from the institute and other universities, we would have maybe never reached the same level of maturity and scientific advancement. Therefore, it is important to thank all colleagues, collaborators, and Master's and Ph.D. students for their contributions. After more than 20 years of research in artificial intelligence and software engineering, accompanied by a huge number of peer-reviewed publications, supervised Master's and Ph.D. students, and projects successfully carried out, it is difficult to focus only on a few things that have been achieved. Hence, in the following, we discuss the most important contributions, picking up two areas, i.e., *diagnosis and debugging*, and *software and system testing* where a combination of software engineering and artificial intelligence techniques and methodologies turned out to be very much beneficial. Besides research, SEAI is heavily involved in teaching from the undergraduate to the PhD level.

4.10.2 Group Structure

SEAI comprises one full professor (Franz Wotawa), one assistant professor (Alexander Perko), one senior lecturer (Roxane Koitz-Hristov), one senior scientist (Iulia Nica), and a changing number of Master and Ph.D. student, and Post-Docs working on externally funded projects. SEAI has been seen as the core research group at our institute since its beginnings in 2003. During this time, SEAI has carried out more than 30 research projects funded by the EU, FFG, FWF, Christian-Doppler Laboratory Society, and other funding sources. In total, the funding of the group exceeds 11 million Euros and has allowed more than 40 Ph.D. students to work on their theses. The largest projects include Softnet Austria, which was an FFG Comet Centre incorporating more than 15 industrial and academic partners in Austra focusing on bringing current software engineering methods into industrial practice, and the Christian-Doppler Laboratory for Quality Assurance Methodologies for Autonomous Cyber-Physical Systems (QAMCAS) funded by the Christian-Doppler Laboratory Society and our industrial partner AVL. Although many projects involve transferring foundational research into practice, the SEAI has always carried out basic research projects funded by the FWF.

4.10.3 Scientific Highlights

SEAI has been conducting research in software engineering, artificial intelligence, and their combination. In this section, we want to focus on two application areas where SEAI substantially contributes to the body of knowledge, i.e., diagnosis and debugging and software and system testing. Before discussing these application areas' two most recent scientific highlights, we give a broad overview of activities and still open challenges. We refer the interested reader to the referenced literature for more information and details.

Before establishing the institute as a separate unit, research focussed mainly on diagnosis and debugging. Early work was on the automation of debugging Java [144, 143] and VHDL programs [202]. In addition, there was work on improving the computation of model-based diagnosis [101, 184] and providing means for relating different diagnosis and debugging techniques [203]. This work has been extended in the following years, see, e.g., [204, 129, 130, 161].

Moreover, we raise awareness of methodologies like model-based diagnosis to a larger audience [160].

From 2004 to 2010, SEAI extended its portfolio to software testing and diagnosis of hybrid systems. Interesting work includes providing a causal analysis for diagnosis of hybrid automata [105], the application of diagnosis to robotics [104] and robust plan execution [178]. In the area of testing, we started utilizing model-checkers for test suite generation [95, 96, 93] and came up with a survey [97]. In addition, we considered the more classical model-based testing [197, 196] and also provided contributions to test suite reduction[94].

From 2011 to 2020, the SEAI group did a lot of work on diagnosis and debugging. In diagnosis, we investigated the use of abduction [205], did work on dependent failures [194], provided a comparative experiment and analysis of different diagnosis algorithms [155], discussed distributed model-based diagnosis [208], and later focused on abductive diagnosis again [134, 135]. Work on the borderline between diagnosis and debugging, i.e., diagnosis of hardware designs, there are two publications worth mentioning, i.e., fault localization in Verilog programs [158] and digital circuits [77]. During this period, there had been a lot of work on debugging SEAI conducted. This includes the combination of different fault localization methods [112, 111, 113] and the application of constraints as modeling methodology for debugging [207]. In addition, the focus has been on niches like spreadsheet debugging [110, 106, 117, 5, 109, 107, 133] also incorporating new ideas like using smells [131] or fragments [116]. Together with colleagues, we contributed to a survey on automated fault localization [201]. In testing, we initially continued on knowledge transfer, i.e., on bringing model-based testing into practice [172]. Furthermore, we provided contributions for integrating testing and mutations into debugging [156], applied testing to other domains like Service-Oriented Architectures (SOA) [119] and security testing, where we used combinatorial testing [175] and AI planning [56] for generating test suites. There is other work on utilizing formal methods for generating test oracles [165] and the utilization of ontologies for test suite generation [139].

In the last 3 years, from 2021 to 2023, we continued working on quality assurance for autonomous driving, coming up with surveys [159, 211] and also with experimental evaluations comparing different test suite generation techniques [128]. In addition, we empirically compared different algorithms for test suite reduction [120]. Regarding spreadsheet debugging, SEAI provided user studies on explainability [148] and tool over-reliance [147], combined different methods for improved spreadsheet fault localization [149] and continued work on using smells and other metrics [108, 132]. In diagnosis, SEAI provided means for utilizing answer set programming [206] and the use of various fault localization techniques to adaptive systems [212, 126, 213]. Further work worth mentioning are contributions to the research field of swarm intelligence [127, 115].

Tool over-reliance

Spreadsheets are a form of end-user programming characterized by having users with limited or no knowledge of software engineering and its methodologies. Hence, users of spreadsheets usually do not know how to ensure quality and detect and localize faults. Any (semi-) automated tool supporting these activities is of great practical interest. In the FWF-funded project iDEOS, we carried out several user studies focusing on different aspects. In the first user study, the objective was to obtain information about the usefulness of debugging tools for spreadsheets and, in particular, the perceived usefulness. In addition, we wanted to clarify how users are interacting with a debugging tool and to identify issues that may arise. To answer these questions, we carried out a user study considering the tool SmellChecker, which utilizes smells and metrics for identifying potentially faulty cells in spreadsheets, which are marked accordingly. The user study considered three user groups, i.e., one control group where SmellChecker was not used, and two treatment groups both using smell checker. However, in one treatment group (group B), we explicitly mention that SmellChecker does not necessarily detect all faults, and in the other group (group A), it does not.

All treatment groups used the same faulty spreadsheet from the ENRON corpus comprising 9 faults. 4 faults are covered by SmellChecker, and 5 are not. For all groups, we asked their participants to identify and describe all faults occurring in the spreadsheet. We used a post-questionnaire to obtain the opinion on the tool and how the user experience was to identify the faults. Participants of the study were selected using the crowd-sourcing platform Prolific10. Overall, 59 participants successfully completed the task online. The selection of participants for the three groups was performed randomly, focusing on an almost equal number of participants per treatment group. On average, participants needed 11.57 minutes to complete the task. It is worth noting that we paid 7 pounds for each participant who completed the task for compensation.

After carrying out the study and analyzing the outcome, we found out that the mean number of correctly identified faults was highest for treatment group B, where we explicitly informed participants about the possibility that SmellChecker might miss fault candidates. In the second place, we have the control group, and finally, treatment group A, which rarely detected more faults than the tool indicated. The detailed results are in Table 4.1. Interestingly, the control group requires the most time (mean 11.28 minutes), followed by treatment group B (10.86 minutes on average) and finally, treatment group A (mean 6.39 minutes). From these observations, it is obvious that candidates from treatment group A trusted the tool more and followed its advice even if faults were missing. Hence, this is a clear indication of tool over-reliance.

In the post-questionnaire, we also obtained useful results. On a scale from 1 (meaning strongly disagree) to 5 (meaning strongly agree), both treatment group participants would like to use the tool frequently, with a very similar value of 4 on average. They also almost agree that the tool was not unnecessarily

Table 4.1: Mean number and standard deviation of correctly identified faults

Group	Mean	Std. dev.
Control	5.60	2.78
Treatment A	4.05	0.58
Treatment B	6.26	2.12

complex but easy to use. The only question where the treatment group results varied more significantly was on the confidence that the tool helped to locate all faulty cells. In treatment group A, more people agreed (with a value of more than 4 on average), whereas this was not the case for treatment group B, which had a value of less than 3 on average. Hence, again, this shows that participants of treatment group A over-rely on the tool's capabilities. For more information and details, we refer to our recent publication [147].

In summary, we state that tool over-reliance is important and might lead to undesired effects, i.e., missing faults when using a debugging tool. Hence, we have to prevent these negative effects. However, there might be cases where over-reliance is not problematic, i.e., in case the tool improves the outcome when compared to manual work (which was not the case in our study) or when a tool itself always produces correct and complete results. For debugging, the paper indicates that work on improving debugging would indeed improve usability.

Testing autonomous driving

This research has been carried out as part of the Christian-Doppler Laboratory QAMCAS, where SEAI developed ontology-based testing [139]. The underlying idea is to use ontologies to represent an environmental model. From the model comprising static and dynamic knowledge, e.g., road networks or vehicle or pedestrian behavior, we extract an input model that can be used by other techniques like search-based testing (SBT) or combinatorial testing (CT) to compute test suites. Ontology-based testing is novel insofar as it focuses on the generation of test cases only considering the environment of a system but not the system itself. Hence, the approach likely produces interactions between the system under test and its environment that the system can hardly handle. In a recent publication [128], we compared SBT, CT, and random testing considering two automated emergency braking systems (AEBs). We used the time to collision as a test oracle for the experimental study. In Figure 4.12 we give an overview of the method and the automated framework for empirical data collection

Before execution, abstract test cases are converted into concrete test scenarios that are executable in simulation, according to the OpenScenario specifi-

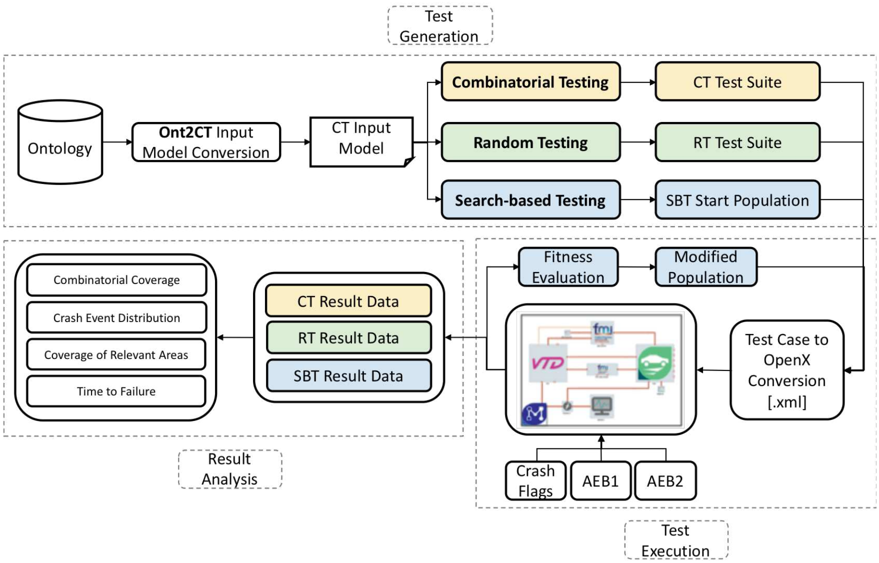

Figure 4.12: Framework for quantitative comparison of search-based and combinatorial testing, considering random testing as the baseline. All three methods share the same CT test input model derived from an AEB ontology to generate and evaluate test suits against two AEB implementations automatically. The collected result data is used for further analysis regarding test coverage and efficiency.

cation[4]. The test execution framework comprises both systems under test (i.e., AEB1 and AEB2), crash flags, and co-simulation, which includes environment and traffic simulation, vehicle dynamics, and sensor models. In detail, we use Vires VTD [5] as a virtual driving environment platform, AVL VSM for vehicle dynamic simulation and AVL Model.CONNECT as the co-simulation tool to connect all the software with the AEB function under test.

During the execution of concrete test scenarios, we monitor time-to-collision (TTC), where the minimum TTC reveals how close the scenario came to a crash situation. A scenario that results in a TTC close to zero seconds indicates a crash (i.e., AEB failure), where we distinguish five different crash events, as shown in Table 4.2.

Table 4.2: Defined crash flags

Crash flags	Description
FCV	Collision with front vehicle
FCP1	Front Collision with pedestrian 1
FCP2	Front Collision with pedestrian 2
SCP1	Side Collision with pedestrian 1
SCP2	Side Collision with pedestrian 2

Here, FCV refers to a crash where the Ego vehicle strikes the rear part of the leading vehicle. FCP1 and SCP1 refer to a front, respectively, side collision with a pedestrian number one, depending on which side of the pedestrian's surrounding bounding box the Ego vehicle strikes. The same applies to FCP1 and SCP2 for pedestrian number two.

In Table 4.3, we summarize the evaluation result for each test generation technique applied to AEB1. If, during scenario execution, a crash event is observed, we consider the scenario as failed, otherwise passed. In the presented table, $Total$ refers to the number of test cases generated, $Pass$ summarizes the number of successful test cases, and $Fail$ refers to the number of failed test cases. For failing test scenarios, we distinguish between five different crash events, as previously described and shown in Table 4.2. Due to their probabilistic elements, search-based test generation (SBT) and random testing (RT) have each been executed ten times against AEB1 and AEB2, marked from SBT01 to SBT10 and RT01 to RT10, respectively. The combinatorial test generation results with strengths two (i.e., CT2_01) and three (i.e., CT3_01) are also presented in the table.

As we observe from the summary table, FCV is the most difficult crash type to detect for all three methods. SBT generated five FCV scenarios in two test runs (SBT05 and SBT10) and RT in just one test run (RT03) out of ten. CT2 test scenarios did not cover an FCV crash event. However, CT3 detected two FCV crash events. The remaining four crash events are detected by each test generation method, some of the less frequent as FCP2, which SBT only identifies in three out of ten test runs (SBT02, SBT05, and SBT06) or SCP2,

[4] Please refer to `http://www.openscenario.org/`

[5] Please refer to `https://vires.com/vtd-vires-virtual-test-drive/`

Table 4.3: SBT, RT, and CT test generation and crash summary for AEB1

SBT	AEB1	Total	Pass	Fail	FCV	FCP1	SCP1	FCP2	SCP2
1	SBT01	1024	806	218	0	162	56	0	1
2	SBT02	983	857	126	0	31	94	1	0
3	SBT03	1002	802	200	0	120	79	0	1
4	SBT04	1010	743	267	0	216	50	0	1
5	SBT05	1005	807	198	3	151	41	4	1
6	SBT06	988	884	104	0	43	55	4	3
7	SBT07	971	743	228	0	203	24	0	5
8	SBT08	990	771	219	0	156	61	0	2
9	SBT09	963	764	199	0	165	34	0	0
10	SBT10	1021	801	220	2	134	82	0	2
RT	AEB1	Total	Pass	Fail	FCV	FCP1	SCP1	FCP2	SCP2
1	RT01	978	949	29	0	19	9	1	0
2	RT02	978	936	42	0	24	17	1	0
3	RT03	978	943	35	1	16	17	1	0
4	RT04	978	938	40	0	24	16	0	0
5	RT05	978	944	34	0	17	16	0	1
6	RT06	978	926	52	0	34	18	0	0
7	RT07	978	949	29	0	12	17	0	0
8	RT08	978	939	39	0	19	19	0	1
9	RT09	978	932	46	0	24	20	2	0
10	RT10	978	937	41	0	21	17	3	0
CT2	AEB1	Total	Pass	Fail	FCV	FCP1	SCP1	FCP2	SCP2
1	CT2_01	978	942	36	0	12	18	4	2
CT3	AEB1	Total	Pass	Fail	FCV	FCP1	SCP1	FCP2	SCP2
1	CT3_01	21418	19567	1851	2	604	999	159	87

which RT only covers in two out of ten test runs (RT05 and RT08).

Furthermore, we observe that each technique can generate a substantial number of scenarios that result in FCP1 or SCP1 events, indicating that AEB1 has noticeable defects in avoiding collision with pedestrian 1.

On average, SBT can generate approximately five times the number of failed test cases compared to RT or CT2, where the average number of failed test cases generated by RT is similar to CT2. CT3 seems to be the most reliable technique concerning crash type coverage, as all crash types are guaranteed to be covered by its test suite. However, the CT3 test suite is twice as large as all executed SBT or RT test runs combined. Further, SBT and RT have also demonstrated that they can eventually detect all crash types, even if not in every test run. Note that for the AEB2 functions, we gain similar results when comparing the different test suite generation methods.

In summary, we showed that ontologies can be used successfully to generate failure-inducing test suites. In addition, we can conclude that combinatorial testing provides the best outcome for fault detection compared to search-based and random testing in automated driving functions and autonomous driving.

4.10.4 Impact Highlights

Besides introducing more than 40 students to foundational and applied research and finally becoming a Ph.D., the SEAI group also has a lot of successful alumni. For example, Gordon Fraser, who finalized his Ph.D. in 2009, is now a full professor at the University of Passau, Germany, Mihai Nica (Ph.D. in 2010) is currently the head of autonomous driving at AVL, Martin Weiglhofer (Ph.D. in 2009) is managing director at Jungheinrich in Austria, Stefan Galler (Ph.D. in 2011) is Vice President of Sales and Business Development at Frequentis USA, and there are more. In addition, to enable successful careers in academia and industry, the SEAI group is also proud to have members receiving prestigious awards like the ACM SIGSOFT Distinguished Paper Award (NIER Track) for the publication “Combining Spreadsheet Smells for Improved Fault Prediction”, published in the Proceedings of the 40th International Conference on Software Engineering, NIER – New Ideas and Emerging Results Track, May 27th-June 3rd, 2018. In Eric Wong, Nikolaos Mittas, Elvira Maria Arvanitou, and Yihao Li, A bibliometric assessment of software engineering themes, scholars and institutions (2013-2020), The Journal of Systems & Software, Vol 180 (2021), https://doi.org/10.1016/j.jss.2021.111029, (Table 9), Franz Wotawa was mentioned on place 9 of the most impactful software engineering researchers, showing clearly the impact of the SEAI group in research. In addition, our partner AVL at the CD laboratory QAMCAS has products and services, e.g., ontology-based test case generation, model-based reasoning for perception systems, and others currently integrated into the portfolio and have already led to substantial revenues. Hence, from the economic perspective, research carried out at the Institute of Software Technology, particularly the SEAI group, is a success story.

Chapter 5

Statistics and Figures

The Institute of Software Technology (IST) has always been active in teaching and research. In the following, we outline some statistics regarding teaching and research activities and relate them to other institutes for benchmarking. At the beginning of December 2023, the IST comprises 21 full-time equivalent researchers funded directly by the university and 3 full-time equivalent administrative staff members. The researchers include 3 full professors, 4 associate professors, 2 holders of a tenure track position who will join the group of associate professors next year, 10 assistant professors, and 2 senior lecturers. The senior lecturers teach at the local universities of applied sciences as part of a common project with our university and are not available to teach at TU Graz. Moreover, some of the 10 assistant professors are also funded via additional projects and are not part of the fixed positions associated with the IST. The following numbers are from an internal document of TU Graz, which is published every year for comparing institutes and faculties considering resources and their use[1]. On average, the number of researchers funded by TU Graz over 2 years was 20.2.

Teaching: On average, over 3 years, the IST taught 164 hours a week. In the study year 2021/22, the members conducted 4,703 exams in total. When comparing these figures with the ones of the other institutes, the IST is in second place, considering the teaching hours and the number of exams. This holds also for the exams when considering the number of available teaching personnel. Only for the semester hours the IST would be in third place. However, interestingly, in this case, the institute holding the first place would be ranked 7th among 10 institutes.

The IST is also very active in supervising Master's and Ph.D. theses. On average, from 2001 to 2023, the IST supervised 14.1 Master's theses and 3.6

[1] Inputbalance & Ergebnissäulen Informatik und Biomedizin. Technik (7000), Beobachtungszeitraum 2022 (Studienjahr 2021/22), Stand Juli 2023, Manuela Berner.

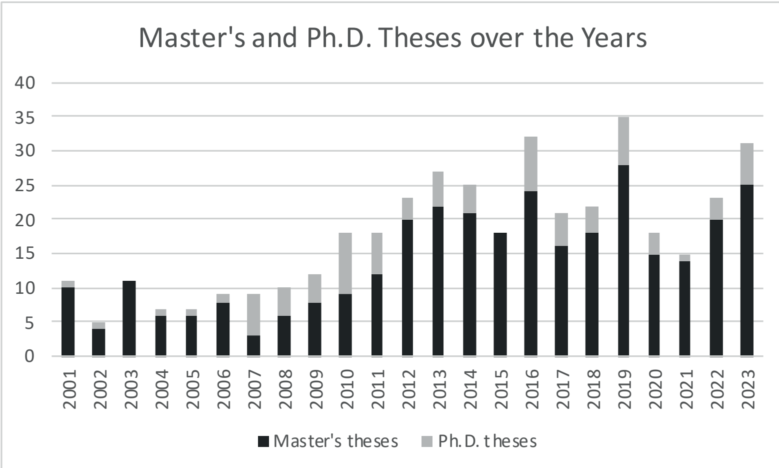

Figure 5.1: Master's and Ph.D. theses supervision from 2001 to 2023

Ph.D. theses per annum. In Figure 5.1, we outline the distribution of supervision over the years, comprising a lot of fluctuations with several peaks. The total number of Master's theses in this period is 324, and for Ph.D. theses, 83. Hence, there is about 1 Ph.D. for 4 Master's students who finalize their studies. It is also interesting to know that several students who finalized a Master's thesis also continue working on a Ph.D. For Franz Wotawa, who supervised 90 Master's students in total at the IST, there are 23 who started a Ph.D. at the IST (but not necessarily under his supervision). Hence, about 1 in 5 Master's students also continue to work in research. The other 17 Ph.D. students finalizing under the supervision of Franz Wotawa were mainly external students from other universities.

Research: There are several metrics someone can use to measure research activities, e.g., the number of papers, the quality of papers considering their publication media, the amount of funding achieved, or the number of researchers not paid by the university directly. In 2022, the IST received 2.166 Million Euros in funding (place 3 of 10). From this funding, 1.557 Million Euros were relevant for the budget in 2022 (place 1 of 10). With this funding, a full-time equivalent of 29.9 Post-Docs and Ph.D. students could be paid for the project and research activities. Again, This is placed 3rd in the ranking of all 10 Faculty of Computer Science and Biomedical Engineering institutes.

Regarding publications, the IST provided 2022 (according to the report "Inputbalance & Ergebnissäulen Informatik und Biomedizin. Technik") 109 papers, not distinguishing conferences from journals, which is about 5.40 papers per researcher funded by the university. For the total number of publications,

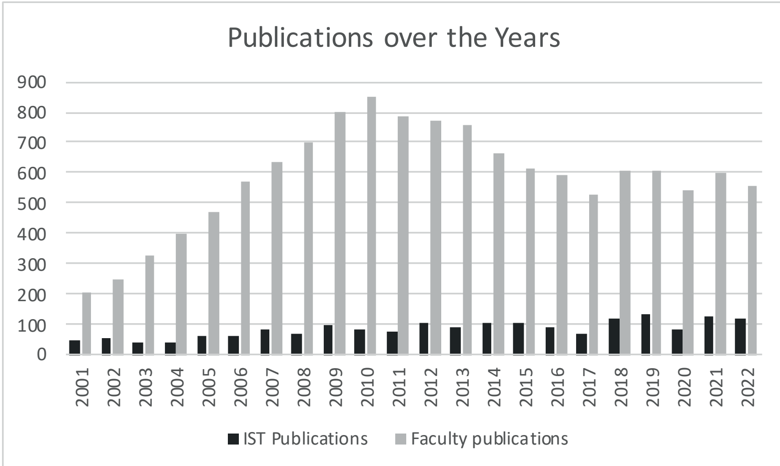

Figure 5.2: Number of IST publications from 2001 to 2023

this is place 2, and for the one considering the available human resources, it is place 1. Hence, in the internal competition, the IST performed very well with respect to the publication activities. It is worth noting that the number of publications of the institute for 2022 is higher in the Pure[2] information system. This might be the case because some publications may be added later and are not considered in the faculty report. However, the IST published about 118 papers per annum over the last five years on average. In Figure 5.2, we depict the number of publications from 2001. With one exception, the number of publications has become stable above 100 in the past 5 years. Interestingly, there was a peak in 2010 of publications from the faculty, which now seems to stabilize at a lower level.

Of course, it is not the number of publications that make a difference in research. The quality has to be ensured, too. Here, we can report that in 2023, we currently have 17, and in 2022, even 28 journal publications. There are many publications in top journals like the Journal of Systems and Software, Machine Learning, Artificial Intelligence, IEEE Transactions on Software Engineering, ACM Computing Surveys, and PloS ONE. Hence, the quality of publications is also given and a major concern.

The Institute of Software Technology has been very successful over the years. In most of the indicators and metrics, the institute ranges within the top 3 of all faculty institutes. With 20.2 full-time employees, the institute holds a share of 15.2% over available staff from direct university funding. However, in all categories, the institute has a higher share, i.e., 16.0% of total funding, 22.7

[2]See `https://graz.elsevierpure.com/en/organisations/institute-of-software-technology-7160`; last accessed December 3rd, 2023.

of total funding that is relevant for the budget, 21.1% of all faculty's publications, and 20.5% of all researchers funded externally. In teaching, the situation is the same, i.e., 18.1% of all lecture hours and 19.4% of all exams. Hence, the Institute always delivers above expectations.

Chapter 6

Conclusions and Future Activities

After reading the book, it becomes apparent that the institute is heterogeneous concerning research topics and activities. The topics range from core areas of software engineering, like software engineering processes and quality assurance, to artificial intelligence, focusing on knowledge-based reasoning, to theoretical computer science, from computational geometry to complexity theory. This heterogeneity is also visible in the kind of research carried out by the different groups. Some are more devoted to curiosity-driven research, having mainly projects funded by the Austrian Science Fund (FWF), to very applied research funded by the European Union and the FFG. What unifies all these activities is the joy of conducting research, the interest in identifying challenges and finding solutions, and educating students to become the next generation of excellent researchers. Hence, all Institute of Software Technology members love science and focus on excellence. In addition, diversity is an enabler for new research activities and projects. There have always been interactions among the groups leading to joint publications. However, there is room for improvement; hopefully, we will see more such interactions among the research groups soon. Celebrating 20 years together, having talks from all areas should help to understand the important topics in the different fields.

Within the next ten years, we will see several changes. Due to retirement, there will be changes in the institute's structure. Successors should be hired according to their excellent standing in the Institute of Software Technology's scientific areas. With new people, new topics will arise. However, with wise selections, the institute will improve and always be at the forefront of research in software engineering, artificial intelligence, and theoretical computer science.

Index

Bibliography

[1] Bernardo M. Ábrego, Oswin Aichholzer, Silvia Fernández-Merchant, Dan McQuillan, Bojan Mohar, Petra Mutzel, Pedro Ramos, R. Bruce Richter, and Birgit Vogtenhuber. Bishellable drawings of K_n. *SIAM Journal on Discrete Mathematics*, 32(4):2482–2492, 2018.

[2] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, T. Hackl, J. Pammer, A. Pilz, P. Ramos, G. Salazar, and B. Vogtenhuber. All good drawings of small complete graphs. In *Proc.* 31^{st} *European Workshop on Computational Geometry EuroCG '15*, pages 57–60, Ljubljana, Slovenia, 2015.

[3] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, P. Ramos, and G. Salazar. More on the crossing number of K_n: Monotone drawings. *Electronic Notes in Discrete Mathematics*, 44:411–414, 2013. Special issue dedicated to LAGOS2013.

[4] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, P. Ramos, and G. Salazar. Shellable drawings and the cylindrical crossing number of K_n. *Discrete and Computational Geometry*, 52:743–753, 2014.

[5] Rui Abreu, Birgit Hofer, Alexandre Perez, and Franz Wotawa. Using constraints to diagnose faulty spreadsheets. *Softw. Qual. J.*, 23(2):297–322, 2015.

[6] Bernhard Aichernig, Harald Brandl, Elisabeth Jöbstl, Willibald Krenn, Rupert Schlick, and Stefan Tiran. MoMuT::UML – model-based mutation testing for UML. In *8th IEEE International Conference on Software Testing, Verification and Validation, ICST 2015, Graz, Austria, April 13-17, 2015*, pages 1–8. IEEE, 2015.

[7] Bernhard K. Aichernig. Mutation Testing in the Refinement Calculus. *Formal Aspects of Computing*, 15(2-3):280–295, 2003.

[8] Bernhard K. Aichernig, Jakob Auer, Elisabeth Jöbstl, Robert Korošec, Willibald Krenn, Rupert Schlick, and Birgit Vera Schmidt. Model-based mutation testing of an industrial measurement device. In *Tests and Proofs - 8th International Conference, TAP 2014, Held as Part of STAF 2014,*

York, UK, July 24-25, 2014. Proceedings, volume 8570 of *Lecture Notes in Computer Science*, pages 1–9. Springer-Verlag, 2014.

[9] Bernhard K. Aichernig, Priska Bauerstätter, Elisabeth Jöbstl, Severin Kann, Robert Korosec, Willibald Krenn, Cristinel Mateis, Rupert Schlick, and Richard Schumi. Learning and statistical model checking of system response times. *Software Quality Journal*, 27(2):757–795, 2019. Open Access.

[10] Bernhard K. Aichernig, Roderick Bloem, Masoud Ebrahimi, Martin Horn, Franz Pernkopf, Wolfgang Roth, Astrid Rupp, Martin Tappler, and Markus Tranninger. Learning a behavior model of hybrid systems through combining model-based testing and machine learning. In Christophe Gaston, Nikolai Kosmatov, and Pascale Le Gall, editors, *Testing Software and Systems - 31st IFIP WG 6.1 International Conference, ICTSS 2019, Paris, France, October 15-17, 2019, Proceedings*, volume 11812 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2019.

[11] Bernhard K. Aichernig, Harald Brandl, Elisabeth Jöbstl, and Willibald Krenn. Model-based mutation testing of hybrid systems. In Frank S. de Boer, Marcello M. Bonsangue, Stefan Hallerstede, and Michael Leuschel, editors, *Formal Methods for Components and Objects - 8th International Symposium, FMCO 2009, Eindhoven, The Netherlands, November 4-6, 2009. Revised Selected Papers*, volume 6286 of *Lecture Notes in Computer Science*, pages 228–249. Springer-Verlag, 2010.

[12] Bernhard K. Aichernig, Harald Brandl, Elisabeth Jöbstl, Willibald Krenn, Rupert Schlick, and Stefan Tiran. Killing strategies for model-based mutation testing. *Journal of Software Testing, Verification and Reliability (STVR)*, 25(8):716–748, 2015.

[13] Bernhard K. Aichernig, Brigitte Fröhlich, and Andreas Kerschbaumer, editors. *Formal Aspects of Software Engineering. Colloquium in Honor of Peter Lucas, May 18–19, 2001*, volume 7, number 8 of *Journal of Universal Computer Science (special issue)*. Springer-Verlag and Know Center Graz, 2001.

[14] Bernhard K. Aichernig and Jifeng He. Mutation testing in UTP. *Formal Aspects of Computing*, 21(1-2):33–64, February 2009.

[15] Bernhard K. Aichernig, Klaus Hörmaier, Florian Lorber, Dejan Nickovic, and Stefan Tiran. Require, test, and trace IT. *International Journal on Software Tools for Technology Transfer (STTT)*, 19:409–426, 2017. Open Access. Published online: 29 November 2016.

[16] Bernhard K. Aichernig, Elisabeth Jöbstl, and Stefan Tiran. Model-based mutation testing via symbolic refinement checking. *Science of Computer Programming*, 97:383–404, 2015.

[17] Bernhard K. Aichernig, Florian Lorber, and Dejan Nickovic. Time for mutants — model-based mutation testing with timed automata. In *Proceedings of the 7th International Conference on Tests and Proofs (TAP 2013), Budapest, Hungary, June 16-20, 2013*, volume 7942 of *Lecture Notes in Computer Science*, pages 20–38. Springer-Verlag, 2013.

[18] Bernhard K. Aichernig, Edi Muskardin, and Andrea Pferscher. Learning-based fuzzing of iot message brokers. In *14th IEEE Conference on Software Testing, Verification and Validation, ICST 2021, Porto de Galinhas, Brazil, April 12-16, 2021*, pages 47–58. IEEE, 2021.

[19] Bernhard K. Aichernig, Andrea Pferscher, and Martin Tappler. From passive to active: Learning timed automata efficiently. In Ritchie Lee, Susmit Jha, and Anastasia Mavridou, editors, *NASA Formal Methods - 12th International Symposium, NFM 2020, Moffett Field, CA, USA, May 11-15, 2020, Proceedings*, volume 12229 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2020.

[20] Bernhard K. Aichernig and Percy Antonio Pari Salas. Test case generation by OCL mutation and constraint solving. In Kai-Yuan Cai and Atsushi Ohnishi, editors, *QSIC 2OO5, Fifth International Conference on Quality Software, Melbourne, Australia, September 19-21, 2005*, pages 64–71. IEEE Computer Society, 2005.

[21] Bernhard K. Aichernig and Martin Tappler. Efficient active automata learning via mutation testing. *Journal of Automated Reasoning*, 63(4):1103–1134, 2019. Open Access.

[22] Bernhard K. Aichernig and Martin Tappler. Probabilistic black-box reachability checking (extended version). *Formal Methods in System Design*, 54(3):416–448, 2019. Open Access.

[23] O. Aichholzer, L.S. Alboul, and F. Hurtado. On flips in polyhedral surfaces. *International Journal of Foundations of Computer Science (IJFCS), special issue on Volume and Surface Triangulations*, 13(2):303–311, 2002.

[24] O. Aichholzer, F. Aurenhammer, and H. Krasser. Enumerating order types for small point sets with applications. *Order*, 19:265–281, 2002.

[25] O. Aichholzer, M. Balko, T. Hackl, J. Kynčl, I. Parada, M. Scheucher, P. Valtr, and B. Vogtenhuber. A superlinear lower bound on the number of 5-holes. *Journal of Combinatorial Theory A*, pages 1–31, 2019. online.

[26] O. Aichholzer, J. García, D. Orden, and P.A. Ramos. New lower bounds for the number of $(\leq k)$-edges and the rectilinear crossing number of K_n. *Discrete & Computational Geometry*, 38:1–14, 2007.

[27] O. Aichholzer, F. Hurtado, and M. Noy. A lower bound on the number of triangulations of planar point sets. *Computational Geometry: Theory and Applications*, 29(2):135–145, 2004.

[28] O. Aichholzer and H. Krasser. Abstract order type extension and new results on the rectilinear crossing number. *Computational Geometry: Theory and Applications, Special Issue on the 21st European Workshop on Computational Geometry*, 36(1):2–15, 2006.

[29] Oswin Aichholzer, Jean Cardinal, Tony Huynh, Kolja Knauer, Torsten Mütze, Raphael Steiner, and Birgit Vogtenhuber. Flip distances between graph orientations. *Algorithmica*, pages 1–28, 2020.

[30] Oswin Aichholzer, Erik D. Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Zuzana Masárová, Mikhail Rudoy, Virginia Vassilevska Williams, and Nicole Wein. Hardness of Token Swapping on Trees. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:15, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[31] Oswin Aichholzer, Frank Duque, Ruy Fabila-Monroy, Carlos Hidalgo-Toscano, and Oscar E. García-Quintero. An Ongoing Project to Improve the Rectilinear and the Pseudolinear Crossing Constants. *Journal of Graph Algorithms and Applications*, 24(3):421–432, 2020.

[32] Oswin Aichholzer, Alfredo García, Javier Tejel, Birgit Vogtenhuber, and Alexandra Weinberger. Twisted Ways to Find Plane Structures in Simple Drawings of Complete Graphs. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry (SoCG 2022)*, volume 224 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:18, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[33] Oswin Aichholzer, Jan Kynčl, Manfred Scheucher, Birgit Vogtenhuber, and Pavel Valtr. On crossing-families in planar point sets. *Computational Geometry*, 107:101–899, 2022.

[34] Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip distance between triangulations of a simple polygon is NP-complete. *Discrete Comput. Geom.*, 54(2):368–389, 2015.

[35] Latifa Al-Abdulkarim, Katie Atkinson, and Trevor J. M. Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.

[36] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[37] Qamar Alfalouji, Thomas Schranz, Basak Falay, Sandra Wilfling, Johannes Exenberger, Thorsten Mattausch, Claudio Gomes, and Gerald Schweiger. Co-simulation for buildings and smart energy systems—a taxonomic review. *Simulation Modelling Practice and Theory*, page 102770, 2023.

[38] Qamar Alfalouji, Thomas Schranz, Alexander Kümpel, Markus Schraven, Thomas Storek, Stephan Gross, Antonello Monti, Dirk Müller, and Gerald Schweiger. Iot middleware platforms for smart energy systems: An empirical expert survey. *Buildings*, 12(5):526, 2022.

[39] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Ricardo Simari, Matthias Thimm, and Serena Villata. Towards artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.

[40] Fahiem Bacchus, Matti Järvisalo, and Ruben Martins. Maximum Satisfiabiliy. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 929–991. IOS Press, 2021.

[41] Werner Bailer, Hermann Fürntratt, Manfred Klopschitz, Dominik Lampel, Roman Lesjak, Matthias Rüther, Susanne Schweitzer, and Gerald Steinbauer-Wagner. Passive guidance for offroad-vehicles. In *International Conference on Robotics in Alpe-Adria Danube Region*, pages 460–468. Springer, 2022.

[42] Alkida Balliu, Sebastian Brandt, Manuela Fischer, Rustam Latypov, Yannic Maus, Dennis Olivetti, and Jara Uitto. Exponential speedup over locality in MPC with optimal memory. In Christian Scheideler, editor, *36th International Symposium on Distributed Computing, DISC 2022, October 25-27, 2022, Augusta, Georgia, USA*, volume 246 of *LIPIcs*, pages 9:1–9:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[43] Alkida Balliu, Keren Censor-Hillel, Yannic Maus, Dennis Olivetti, and Jukka Suomela. Locally checkable labelings with small messages. In Seth Gilbert, editor, *35th International Symposium on Distributed Computing, DISC 2021, October 4-8, 2021, Freiburg, Germany (Virtual Conference)*, volume 209 of *LIPIcs*, pages 8:1–8:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[44] Alkida Balliu, Rustam Latypov, Yannic Maus, Dennis Olivetti, and Jara Uitto. Optimal deterministic massively parallel connectivity on forests. In

Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2589–2631. SIAM, 2023.

[45] Philipp Bamberger, Fabian Kuhn, and Yannic Maus. Efficient deterministic distributed coloring with small bandwidth. In *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 243–252, 2020.

[46] Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, 2018.

[47] Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo Ricardo Simari, editors, *Proc. COMMA*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 75–86. IOS Press, 2010.

[48] Ringo Baumann, Sylvie Doutre, Jean-Guy Mailly, and Johannes P. Wallner. Enforcement in formal argumentation. In Dov Gabbay, Massimiliano Giacomin, Guillermo R. Simari, and Matthias Thimm, editors, *Handbook of Formal Argumentation*, volume 2, chapter 8, pages 445–510. College Publications, 2021.

[49] Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[50] Theresa Boiger and Gerald Schweiger. Shp2sim: a python pipeline for modelica based district and urban scale energy simulations. *International Journal of Sustainable Energy*, 42(1):1028–1041, 2023.

[51] Andrei Bondarenko, Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93:63–101, 1997.

[52] Marco De Bortoli, Lukáš Chrpa, Martin Gebser, and Gerald Steinbauer-Wagner. Enhancing temporal planning by sequential macro-actions. In *European Conference on Logics in Artificial Intelligence (JELIA 2023)*, number 14281 in Lecture Notes in Computer Science, pages 595–604. Springer, 2023.

[53] Marco De Bortoli and Gerald Steinbauer-Wagner. Evaluating action-based temporal planners performance in the robocup logistics league. In *RoboCup 2022: - Robot World Cup XXV*, number 13561 in Lecture Notes in Computer Science, 2023.

[54] Marco De Bortoli and Gerald Steinbauer-Wagner. Improving applicability of planning in the robocup logistics league using macro-actions refinement. In *RoboCup 2023: Robot World Cup XXVI*, number 14140 in Lecture Notes in Computer Science. Springer, 2024.

[55] Marco De Bortoli and Gerald Steinbauer-Wagner. Temporal planning and acting in dynamic domains. In *RoboCup 2023: Robot World Cup XXVI*, number 14140 in Lecture Notes in Computer Science. Springer, 2024.

[56] Josip Bozic and Franz Wotawa. Planning-based security testing of web applications with attack grammars. *Softw. Qual. J.*, 28(1):307–334, 2020.

[57] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[58] Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczynski. Answer set programming at a glance. *Communications of the ACM*, 54(12):92–103, 2011.

[59] Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes P. Wallner, and Stefan Woltran. Abstract dialectical frameworks. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, chapter 5, pages 237–285. College Publications, 2018.

[60] Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5-6):286–310, 2007.

[61] Zhelun Chen, Zheng O'Neill, Jin Wen, Ojas Pradhan, Tao Yang, Xing Lu, Guanjing Lin, Shohei Miyata, Seungjae Lee, Chou Shen, et al. A review of data-driven fault detection and diagnostics for building hvac systems. *Applied Energy*, 339:121030, 2023.

[62] Kristijonas Cyras, Antonio Rago, Emanuele Albini, Pietro Baroni, and Francesca Toni. Argumentative XAI: A survey. In Zhi-Hua Zhou, editor, *Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 4392–4399. ijcai.org, 2021.

[63] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

[64] Hamid Didari, Eva Reitbauer, Christoph Schmied, Matthias Eder, Stephan Schraml, Nizar Aourik, Marius Feilhauer, Rene Kastner, Roland Hieslmayr, and Gerald Steinbauer-Wagner. Robo-mole: An assistant robot system for first responders in subterranean scenarios. In *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 99–106, 2023.

[65] Hamid Didari and Gerald Steinbauer-Wagner. Lidar-based scene understanding for navigation in unstructured environments. In *Advances in Service and Industrial Robotics*. Springer Nature Switzerland, 2023.

[66] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L Vrabie, et al. All you need to know about model predictive control for buildings. *Annual Reviews in Control*, 50:190–232, 2020.

[67] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.

[68] Johann Eck, Sabine Hirschmugl-Gaisch, Martin Kandlhofer, and Gerald Steinbauer. A cross-generational robotics project day: pre-school children, pupils and grandparents learn together. *Journal of Automation Mobile Robotics and Intelligent Systems*, 8(1), 2014.

[69] Thomas Eckstein, Marco De Bortoli, and Gerald Steinbauer-Wagner. Fixing pddl domains using delta debugging. In *34th International Workshop on Principles of Diagnosis*, Loma Mar, CA, USA, 2023.

[70] Matthias Eder, Raphael Prinz, Florian Schöggl, and Gerald Steinbauer-Wagner. Traversability analysis for off-road environments using locomotion experiments and earth observation data. *Robotics and Autonomous Systems*, 168, 2023.

[71] Matthias Eder, Michael Reip, and Gerald Steinbauer. Creating a robot localization monitor using particle filter and machine learning approaches. *Applied Intelligence*, 52(6), 2022.

[72] Matthias Eder and Gerald Steinbauer-Wagner. A fast method for explanations of failures in optimization- based robot motion planning. In *Advances in Service and Industrial Robotics: RAAD 2022*, volume 120 of *Mechanisms and Machine Science*. Springer, 2022.

[73] Matthias Eder and Gerald Steinbauer-Wagner. Predicting energy consumption and traversal time of ground robots for outdoor navigation on multiple types of terrain. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023.

[74] Matthias Eder and Gerald Steinbauer-Wagner. Robot-dependent traversability estimation for outdoor environments using deep multimodal variational autoencoders. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.

[75] Matthias Eder and Gerald Steinbauer-Wagner. Why did my robot choose this path? explainable path planning for off-road navigation. In *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, pages 139–145. IEEE, 2024.

[76] Basak Falay, Sandra Wilfling, Qamar Alfalouji, Johannes Exenberger, Thomas Schranz, Christian Møldrup Legaard, Ingo Leusbrock, and Gerald Schweiger. Coupling physical and machine learning models: Case study of a single-family house. In *Modelica Conferences*, pages 335–341, 2021.

[77] Alexander Feldman, Ingo Pill, Franz Wotawa, Ion Matei, and Johan de Kleer. Efficient model-based diagnosis of sequential circuits. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2814–2821. AAAI Press, 2020.

[78] Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalcic. *Group Recommender Systems: An Introduction*. Springer Publishing Company Inc., 2nd edition, 2023.

[79] Alexander Felfernig and Robin Burke. Constraint-Based Recommender Systems: Technologies and Research Issues. In *10th Intl. Conference on Electronic Commerce*, ICEC '08, New York, NY, USA, 2008. Association for Computing Machinery.

[80] Alexander Felfernig, Gerhard Friedrich, Bartosz Gula, Martin Hitz, Thomas Kruggel, Gerhard Leitner, Rudolf Melcher, Daniela Riepan, Sabine Strauss, Erich Teppan, and Oliver Vitouch. Persuasive recommendation: Serial position effects in knowledge-based recommender systems. In *Proceedings of the 2nd International Conference on Persuasive Technology*, PERSUASIVE'07, page 283–294, Berlin, Heidelberg, 2007. Springer-Verlag.

[81] Alexander Felfernig, Gerhard Friedrich, Dietmar Jannach, and Markus Stumptner. Consistency-based Diagnosis of Configuration Knowledge Bases. *Artificial Intelligence*, 152(2):213–234, 2004.

[82] Alexander Felfernig, Gerhard Friedrich, Dietmar Jannach, and Markus Zanker. An Integrated Environment for the Development of Knowledge-based Recommender Applications. *Intl. Journal of Electronic Commerce (IJEC)*, 11(2):11–34, 2006.

[83] Alexander Felfernig, Gerhard Friedrich, Monika Schubert, Monika Mandl, Markus Mairitsch, and Erich Teppan. Plausible Repairs for Inconsistent Requirements. In *21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, pages 791–796, Pasadena, California, USA, 2009.

[84] Alexander Felfernig, Thomas Gruber, Guenter Brandner, Paul Blazek, and Martin Stettinger. Customizing Events with EventHelpr. In *8th International Conference on Mass Customization and Personalization –*

Community of Europe (MCP-CE 2018), pages 88–91, Novi Sad, Serbia, 2018.

[85] Alexander Felfernig, Lother Hotz, Claire Bagley, and Juha Tiihonen. *Knowledge-based Configuration: From Research to Business Cases*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2014.

[86] Alexander Felfernig, Klaus Isak, Kalman Szabo, and Peter Zachar. The vita financial services sales support environment. In *Proceedings of the 19th National Conference on Innovative Applications of Artificial Intelligence - Volume 2*, IAAI'07, page 1692–1699. AAAI Press, 2007.

[87] Alexander Felfernig, Monika Schubert, and Stefan Reiterer. Personalized diagnosis for over-constrained problems. In *23rd International Joint Conference on Artificial Intelligence*, IJCAI '13, page 1990–1996. AAAI Press, 2013.

[88] Alexander Felfernig, Monika Schubert, and Christoph Zehentner. An Efficient Diagnosis Algorithm for Inconsistent Constraint Sets. *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing (AIEDAM)*, 26(1):53–62, 2012.

[89] Alexander Felfernig, Martin Stettinger, Müslüm Atas, Ralph Samer, Jennifer Nerlich, Simon Scholz, Juha Tiihonen, and Mikko Raatikainen. Towards utility-based prioritization of requirements in open source environments. In *2018 IEEE 26th International Requirements Engineering Conference (RE)*, pages 406–411, 2018.

[90] Alexander Felfernig, Manfred Wundara, Thi Ngoc Trang Tran, Viet-Man Le, Sebastian Lubos, and Seda Polat-Erdeniz. Sports recommender systems: Overview and research issues. In *arXiv: 2312.03785*, 2023.

[91] Alexander Felfernig, Manfred Wundara, Thi Ngoc Trang Tran, Seda Polat-Erdeniz, Sebastian Lubos, Merfat El Mansi, Damian Garber, and Viet-Man Le. Recommender systems for sustainability: overview and research issues. *Frontiers in Big Data*, 6, 2023.

[92] Alexander Felfernig, Christoph Zehentner, Gerald Ninaus, Harald Grabner, Walid Maalej, Dennis Pagano, Leopold Weninger, and Florian Reinfrank. Group decision support for requirements negotiation. In *Proceedings of the 19th International Conference on Advances in User Modeling*, UMAP'11, page 105–116, Berlin, Heidelberg, 2011. Springer-Verlag.

[93] Gordon Fraser, Angelo Gargantini, and Franz Wotawa. On the order of test goals in specification-based testing. *J. Log. Algebraic Methods Program.*, 78(6):472–490, 2009.

[94] Gordon Fraser and Franz Wotawa. Redundancy based test-suite reduction. In Matthew B. Dwyer and Antónia Lopes, editors, *Fundamental Approaches to Software Engineering, 10th International Conference, FASE 2007, Held as Part of the Joint European Conferences, on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4422 of *Lecture Notes in Computer Science*, pages 291–305. Springer, 2007.

[95] Gordon Fraser and Franz Wotawa. Using model-checkers to generate and analyze property relevant test-cases. *Softw. Qual. J.*, 16(2):161–183, 2008.

[96] Gordon Fraser, Franz Wotawa, and Paul Ammann. Issues in using model checkers for test case generation. *J. Syst. Softw.*, 82(9):1403–1418, 2009.

[97] Gordon Fraser, Franz Wotawa, and Paul Ammann. Testing with model checkers: a survey. *Softw. Test. Verification Reliab.*, 19(3):215–261, 2009.

[98] Laurent Frering, Matthias Eder, Bettina Kubicek, Dietrich Albert, Denis Kalkofen, Thomas Gschwandtner, Heimo Krajnz, and Gerald Steinbauer-Wagner. Enabling and assessing trust when cooperating with robots in disaster response (easier). 2022.

[99] Laurent Frering, Armin Koefler, Michael Huber, Sandra Pfister, Richard Feischl, Alexander Almer, and Gerald Steinbauer-Wagner. Multi-robot support system for fighting wildfires in challenging environments: System design and field test report. In *2023 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 32–38. IEEE, 2023.

[100] Mohsen Ghaffari, Fabian Kuhn, and Yannic Maus. On the complexity of local distributed graph problems. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 784–797. ACM, 2017.

[101] Georg Gottlob, Martin Hutle, and Franz Wotawa. Combining hypertree, bicomp, and hinge decomposition. In Frank van Harmelen, editor, *Proceedings of the 15th European Conference on Artificial Intelligence, ECAI'2002, Lyon, France, July 2002*, pages 161–165. IOS Press, 2002.

[102] John Grant and Maria Vanina Martinez, editors. *Measuring Inconsistency in Information*, volume 73 of *Studies in Logic*. College Publications, 2018.

[103] Magnús M. Halldórsson, Fabian Kuhn, Yannic Maus, and Tigran Tonoyan. Efficient randomized distributed coloring in CONGEST. In *STOC '21:*

53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021, pages 1180–1193, 2021. Full version at CoRR abs/2105.04700.

[104] Michael W. Hofbaur, Johannes Köb, Gerald Steinbauer, and Franz Wotawa. Improving robustness of mobile robots using model-based reasoning. *J. Intell. Robotic Syst.*, 48(1):37–54, 2007.

[105] Michael W. Hofbaur and Franz Wotawa. A causal analysis method for concurrent hybrid automata. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pages 840–846. AAAI Press, 2006.

[106] Birgit Hofer, Rui Abreu, Alexandre Perez, and Franz Wotawa. Generation of relevant spreadsheet repair candidates. In Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan, editors, *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 1027–1028. IOS Press, 2014.

[107] Birgit Hofer, Andrea Hofler, and Franz Wotawa. Combining models for improved fault localization in spreadsheets. *IEEE Trans. Reliab.*, 66(1):38–53, 2017.

[108] Birgit Hofer, Dietmar Jannach, Patrick W. Koch, Konstantin Schekotihin, and Franz Wotawa. Product metrics for spreadsheets - A systematic review. *J. Syst. Softw.*, 175:110910, 2021.

[109] Birgit Hofer, Alexandre Perez, Rui Abreu, and Franz Wotawa. On the empirical evaluation of similarity coefficients for spreadsheets fault localization. *Autom. Softw. Eng.*, 22(1):47–74, 2015.

[110] Birgit Hofer, André Riboira, Franz Wotawa, Rui Abreu, and Elisabeth Getzner. On the empirical evaluation of fault localization techniques for spreadsheets. In Vittorio Cortellessa and Dániel Varró, editors, *Fundamental Approaches to Software Engineering - 16th International Conference, FASE 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7793 of *Lecture Notes in Computer Science*, pages 68–82. Springer, 2013.

[111] Birgit Hofer and Franz Wotawa. Combining slicing and constraint solving for better debugging: The CONBAS approach. *Adv. Softw. Eng.*, 2012:628571:1–628571:18, 2012.

[112] Birgit Hofer and Franz Wotawa. Spectrum enhanced dynamic slicing for better fault localization. In Luc De Raedt, Christian Bessiere, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter J. F. Lucas, editors, *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 420–425. IOS Press, 2012.

[113] Birgit Hofer, Franz Wotawa, and Rui Abreu. AI for the win: improving spectrum-based fault localization. *ACM SIGSOFT Softw. Eng. Notes*, 37(6):1–8, 2012.

[114] Raphael Hoheneder, Alexander Buchegger, Hamid Didari, Matthias Eder, Konstantin Mautner-Lassnig, Clemens Mühlbacher, and Gerald Steinbauer-Wagner. An event-based approach to autonomous navigation. In *Austrian Robotics Workshop (ARW)*, Linz, Austria, 2023.

[115] Hannes Hornischer, Joshua Cherian Varughese, Ronald Thenius, Franz Wotawa, Manfred Füllsack, and Thomas Schmickl. CIMAX: collective information maximization in robotic swarms using local communication. *Adapt. Behav.*, 29(3), 2021.

[116] Dietmar Jannach, Thomas Schmitz, Birgit Hofer, Konstantin Schekotihin, Patrick W. Koch, and Franz Wotawa. Fragment-based spreadsheet debugging. *Autom. Softw. Eng.*, 26(1):203–239, 2019.

[117] Dietmar Jannach, Thomas Schmitz, Birgit Hofer, and Franz Wotawa. Avoiding, finding and fixing spreadsheet errors - A survey of automated approaches for spreadsheet QA. *J. Syst. Softw.*, 94:129–150, 2014.

[118] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich. *Recommender Systems: An Introduction*. Cambridge University Press, USA, 1st edition, 2010.

[119] Seema Jehan, Ingo Pill, and Franz Wotawa. BPEL integration testing. In Alexander Egyed and Ina Schaefer, editors, *Fundamental Approaches to Software Engineering - 18th International Conference, FASE 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9033 of *Lecture Notes in Computer Science*, pages 69–83. Springer, 2015.

[120] Seema Jehan and Franz Wotawa. An empirical study of greedy test suite minimization techniques using mutation coverage. *IEEE Access*, 11:65427–65442, 2023.

[121] Simon Jusner, Stefan Moser, Sebastian Schaffler-Glößl, Richard Halatschek, Matthias Eder, and Gerald Steinbauer-Wagner. A mission architecture for a human-robot collaborative planetary exploration cascade. In *2024 International Conference on Space Robotics (iSpaRo)*, pages 321–327, 2024.

[122] Martin Kandlhofer and Gerald Steinbauer. Evaluating the impact of educational robotics on pupils' technical- and social-skills and science related attitudes. *Robotics and Autonomous Systems*, 75(B):679 – 685, 2016.

[123] Martin Kandlhofer, Gerald Steinbauer, Julia P. Laßnig, Manuel Menzinger, Wilfried Baumann, Margit Ehardt-Schmiederer, Ronald Bieber, Thomas Winkler, Sandra Plomer, Inge Strobl-Zuchtriegl, Marlene Miglbauer, Aron Ballagi, Claudiu Pozna, Gabor Miltenyi, and Imre Szalay Istvan Alfoldi. Edlris: A european driving license for robots and intelligent systems. *Künstliche Intelligenz*, 35(2), 2021.

[124] Martin Kandlhofer, Gerald Steinbauer, Petra Sundström, and Astrid Weiss. Evaluating the long-term impact of RoboCupJunior: A first investigation. In *International Conference on Robotics in Education*, Prague, Czech Republic, 2012.

[125] H. Karloff, S. Suri, and S. Vassilvitskii. A model of computation for mapreduce. 2010.

[126] David Kaufmann, Iulia Nica, and Franz Wotawa. Intelligent agents diagnostics - enhancing cyber-physical systems with self-diagnostic capabilities. *Adv. Intell. Syst.*, 3(5):2000218, 2021.

[127] Daniela Kengyel, Payam Zahadat, Franz Wotawa, and Thomas Schmickl. Towards swarm level optimisation: the role of different movement patterns in swarm systems. *Int. J. Parallel Emergent Distributed Syst.*, 34(3):241–259, 2019.

[128] Florian Klück, Yihao Li, Jianbo Tao, and Franz Wotawa. An empirical comparison of combinatorial testing and search-based testing in the context of automated and autonomous driving systems. *Inf. Softw. Technol.*, 160:107225, 2023.

[129] Daniel Köb and Franz Wotawa. Introducing alias information into model-based debugging. In Ramón López de Mántaras and Lorenza Saitta, editors, *Proceedings of the 16th Eureopean Conference on Artificial Intelligence, ECAI'2004, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain, August 22-27, 2004*, pages 833–837. IOS Press, 2004.

[130] Daniel Köb and Franz Wotawa. Fundamentals of debugging using a resolution calculus. In Luciano Baresi and Reiko Heckel, editors, *Fundamental Approaches to Software Engineering, 9th International Conference, FASE 2006, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 27-28, 2006, Proceedings*, volume 3922 of *Lecture Notes in Computer Science*, pages 278–292. Springer, 2006.

[131] Patrick W. Koch, Birgit Hofer, and Franz Wotawa. On the refinement of spreadsheet smells by means of structure information. *J. Syst. Softw.*, 147:64–85, 2019.

[132] Patrick W. Koch, Konstantin Schekotihin, Dietmar Jannach, Birgit Hofer, and Franz Wotawa. Metric-based fault prediction for spreadsheets. *IEEE Trans. Software Eng.*, 47(10):2195–2207, 2021.

[133] Patrick W. Koch, Konstantin Schekotihin, Dietmar Jannach, Birgit Hofer, Franz Wotawa, and Thomas Schmitz. Combining spreadsheet smells for improved fault prediction. In Andrea Zisman and Sven Apel, editors, *Proceedings of the 40th International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, pages 25–28. ACM, 2018.

[134] Roxane Koitz and Franz Wotawa. Applying algorithm selection to abductive diagnostic reasoning. *Appl. Intell.*, 48(11):3976–3994, 2018.

[135] Roxane Koitz-Hristov and Franz Wotawa. Faster horn diagnosis - a performance comparison of abductive reasoning algorithms. *Appl. Intell.*, 50(5):1558–1572, 2020.

[136] Viet-Man Le, Alexander Felfernig, Thi Ngoc Trang Tran, and Mathias Uta. INFORMEDQX: Informed Conflict Detection for Over-Constrained Problems. In *38th Annual AAAI Conference on Artificial Intelligence*, page to appear, Vancouver, Canada, 2024. AAAI.

[137] Viet-Man Le, Cristian Vidal Silva, Alexander Felfernig, David Benavides, José Galindo, and Thi Ngoc Trang Tran. Fastdiagp: An algorithm for parallelized direct diagnosis. AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023.

[138] Christian Møldrup Legaard, Daniella Tola, Thomas Schranz, Hugo Daniel Macedo, and Peter Gorm Larsen. A universal mechanism for implementing functional mock-up units. In *SIMULTECH*, pages 121–129, 2021.

[139] Yihao Li, Jianbo Tao, and Franz Wotawa. Ontology-based test generation for automated and autonomous driving functions. *Inf. Softw. Technol.*, 117, 2020.

[140] N. Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.

[141] Thomas Linsbichler, Marco Maratea, Andreas Niskanen, Johannes P. Wallner, and Stefan Woltran. Advanced algorithms for abstract dialectical frameworks based on complexity analysis of subclasses and SAT solving. *Artificial Intelligence*, 307:103697, 2022.

[142] Yannic Maus and Jara Uitto. Efficient CONGEST algorithms for the lovász local lemma. In Seth Gilbert, editor, *35th International Symposium on Distributed Computing, DISC 2021, October 4-8, 2021, Freiburg, Germany (Virtual Conference)*, volume 209 of *LIPIcs*, pages 31:1–31:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[143] Wolfgang Mayer, Markus Stumptner, Dominik Wieland, and Franz Wotawa. Can AI help to improve debugging substantially? debugging experiences with value-based models. In Frank van Harmelen, editor, *Proceedings of the 15th European Conference on Artificial Intelligence, ECAI'2002, Lyon, France, July 2002*, pages 417–421. IOS Press, 2002.

[144] Wolfgang Mayer, Markus Stumptner, Dominik Wieland, and Franz Wotawa. Towards an integrated debugging environment. In Frank van Harmelen, editor, *Proceedings of the 15th European Conference on Artificial Intelligence, ECAI'2002, Lyon, France, July 2002*, pages 422–426. IOS Press, 2002.

[145] Sun Meng, Farhad Arbab, Bernhard K. Aichernig, Lacramioara Astefanoaei, Frank S. de Boer, and Jan Rutten. Connectors as designs: Modeling, refinement and test case generation. *Science of Computer Programming*, 77(7/8):799–822, 2012.

[146] Sanjay Modgil and Henry Prakken. A general account of argumentation with preferences. *Artificial Intelligence*, 195:361–397, 2013.

[147] Adil Mukhtar, Birgit Hofer, Dietmar Jannach, and Franz Wotawa. Spreadsheet debugging: The perils of tool over-reliance. *J. Syst. Softw.*, 184:111119, 2022.

[148] Adil Mukhtar, Birgit Hofer, Dietmar Jannach, and Franz Wotawa. Explaining software fault predictions to spreadsheet users. *J. Syst. Softw.*, 201:111676, 2023.

[149] Adil Mukhtar, Birgit Hofer, Dietmar Jannach, Franz Wotawa, and Konstantin Schekotihin. Boosting spectrum-based fault localization for spreadsheets with product metrics in a learning approach. In *37th IEEE/ACM International Conference on Automated Software Engineering, ASE 2022, Rochester, MI, USA, October 10-14, 2022*, pages 175:1–175:5. ACM, 2022.

[150] Edi Muskardin, Bernhard K. Aichernig, Ingo Pill, Andrea Pferscher, and Martin Tappler. Aalpy: an active automata learning library. *Innovations*

in Systems and Software Engineering, a NASA Journal, 18(3):417–426, 2022. Open Access.

[151] Clemens Mühlbacher, Stephan Gspandl, Michael Reip, and Gerald Steinbauer. Improving dependability of industrial transport robots using model-based techniques. In *IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, 2016.

[152] Clemens Mühlbacher and Gerald Steinbauer. Knowledge-aware execution of programs in indigolog. In *International Workshop on Cognitive Robotics (CogRob)*, Prague, Czech Republic, 2014.

[153] Clemens Mühlbacher and Gerald Steinbauer. Belief management using the action history and consistency-based-diagnosis. In *International Workshop on Principles of Diagnosis (DX)*, Denver, CO, USA, 2016.

[154] Clemens Mühlbacher and Gerald Steinbauer. Hands off - a holistic model-based approach for long- term autonomy. In *ICRA Workshop on AI for Long-Term Autonomy*, Stockholm, Sweden, 2016.

[155] Iulia Nica, Ingo Pill, Thomas Quaritsch, and Franz Wotawa. The route to success - A performance comparison of diagnosis algorithms. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 1039–1045. IJCAI/AAAI, 2013.

[156] Mihai Nica, Simona Nica, and Franz Wotawa. On the use of mutations and testing for debugging. *Softw. Pract. Exp.*, 43(9):1121–1142, 2013.

[157] Daphne Odekerken, Floris Bex, AnneMarie Borg, and Bas Testerink. Approximating stability for applied argument-based inquiry. *Intelligent Systems with Applications*, 16:200110, 2022.

[158] Bernhard Peischl, Naveed Riaz, and Franz Wotawa. Automated debugging of verilog designs. *Int. J. Softw. Eng. Knowl. Eng.*, 22(5):695, 2012.

[159] Bernhard Peischl, Oliver A. Tazl, and Franz Wotawa. Testing anticipatory systems: A systematic mapping study on the state of the art. *J. Syst. Softw.*, 192:111387, 2022.

[160] Bernhard Peischl and Franz Wotawa. Model-based diagnosis or reasoning from first principles. *IEEE Intell. Syst.*, 18(3):32–37, 2003.

[161] Bernhard Peischl and Franz Wotawa. Automated source-level error localization in hardware designs. *IEEE Des. Test Comput.*, 23(1):8–19, 2006.

[162] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.

[163] Yannick Pencol, Gerald Steinbauer, Clemens Mühlbacher, and Louise Trave-Massuyes. Diagnosing discrete event systems using nominal models only. In *International Workshop on Principles of Diagnosis (DX)*, 2017.

[164] Andrea Pferscher and Bernhard K. Aichernig. Stateful black-box fuzzing of bluetooth devices using automata learning. In Jyotirmoy V. Deshmukh, Klaus Havelund, and Ivan Perez, editors, *NASA Formal Methods - 14th International Symposium, NFM 2022, Pasadena, CA, USA, May 24-27, 2022, Proceedings*, volume 13260 of *Lecture Notes in Computer Science*, pages 373–392. Springer, 2022.

[165] Ingo Pill and Franz Wotawa. Automated generation of (F)LTL oracles for testing and debugging. *J. Syst. Softw.*, 139:124–141, 2018.

[166] Andrei Popescu, Seda Polat-Erdeniz, Alexander Felfernig, Mathias Uta, Müslüm Atas, Viet-Man Le, Klaus Pilsl, Martin Enzelsberger, and Thi Ngoc Trang Tran. An overview of machine learning techniques in constraint solving. *Journal of Intelligent Information Systems*, 58(1):91–118, 2022.

[167] Václav Rozhon and Mohsen Ghaffari. Polylogarithmic-time deterministic network decomposition and distributed derandomization. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 350–363. ACM, 2020.

[168] Thomas Schranz, Qamar Alfalouji, Thomas Hirsch, and Gerald Schweiger. An open iot platform: Lessons learned from a district energy system. In *2022 Second International Conference on Sustainable Mobility Applications, Renewables and Technology (SMART)*, pages 1–9. IEEE, 2022.

[169] Thomas Schranz, Johannes Exenberger, Christian Møldrup Legaard, Jan Drgona, and Gerald Schweiger. Energy prediction under changed demand conditions: Robust machine learning models and input feature combinations. In *Building Simulation 2021*, volume 17, pages 3268–3275. IBPSA, 2021.

[170] Thomas Schranz, Christian Møldrup Legaard, Daniella Tola, and Gerald Schweiger. Portable runtime environments for python-based fmus: Adding docker support to unifmu. In *Modelica Conferences*, pages 419–424, 2021.

[171] Thomas Schranz, Gerald Schweiger, Siegfried Pabst, and Franz Wotawa. Machine learning for water supply supervision. In *Trends in Artificial Intelligence Theory and Applications. Artificial Intelligence Practices: 33rd*

International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems, IEA/AIE 2020, Kitakyushu, Japan, September 22-25, 2020, Proceedings 33, pages 238–249. Springer, 2020.

[172] Christian Schwarzl and Franz Wotawa. Test case generation in practice for communicating embedded systems. *Elektrotech. Informationstechnik*, 128(6):240–244, 2011.

[173] Gerald Schweiger, Cláudio Gomes, Georg Engel, Irene Hafner, J Schoeggl, Alfred Posch, and T Nouidui. An empirical survey on co-simulation: Promising standards, challenges and research needs. *Simulation modelling practice and theory*, 95:148–163, 2019.

[174] Jorge Santos Simon, Clemens Mühlbacher, and Gerald Steinbauer. Automatic model generation to diagnose autonomous systems. In *International Workshop on Principles of Diagnosis (DX)*, Paris, France, 2015.

[175] Dimitris E. Simos, Josip Bozic, Bernhard Garn, Manuel Leithner, Feng Duan, Kristoffer Kleine, Yu Lei, and Franz Wotawa. Testing TLS using planning-based combinatorial methods and execution framework. *Softw. Qual. J.*, 27(2):703–729, 2019.

[176] Gerald Steinbauer, Martin Kandlhofer, Tara Chklovski, Fredrik Heintz, and Sven Koenig. A Differentiated Discussion About AI Education K-12. *Künstliche Intelligenz*, 35(2), 2021.

[177] Gerald Steinbauer and Franz Wotawa. Detecting and locating faults in the control software of autonomous mobile robots. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, UK, 2005.

[178] Gerald Steinbauer and Franz Wotawa. Robust plan execution using model-based reasoning. *Adv. Robotics*, 23(10):1315–1326, 2009.

[179] Gspandl Stephan, Pill Ingo Hans, Reip Michael, Steinbauer Gerald, and Ferrein Alexander. Belief Management for High-Level Robot Programs. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Barcelona, Spain, 2011.

[180] Gspandl Stephan, Podesser Siegfried, Reip Michael, Steinbauer Gerald, and Wolfram Mate. A Dependable Perception-Decision-Execution Cycle for Autonomous Robots. In *International Conference on Robotics and Automation (ICRA)*, St. Paul, Minnesota, USA, 2012. IEEE.

[181] Martin Stettinger, Alexander Felfernig, Gerhard Leitner, Stefan Reiterer, and Michael Jeran. Counteracting serial position effects in the choicla group decision support environment. IUI '15, pages 148–157, New York, NY, USA, 2015. Association for Computing Machinery.

[182] Martin Stettinger, Thi Ngoc Trang Tran, Ingo Pribik, Gerhard Leitner, Alexander Felfernig, Ralph Samer, Müslüm Atas, and Manfred Wundara. KNOWLEDGECHECKR: Intelligent Techniques for Counteracting Forgetting. In G. DeGiacomo, A. Catala, B. Dilkina, M. Milano, S. Barro, A. Bugarin, and J. Lang, editors, *24th European Conference on Artificial Intelligence (ECAI 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 3034–3039, Santiago de Compostela, Spain, 2020. IOS Press.

[183] Hannes Strass and Johannes P. Wallner. Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory. *Artificial Intelligence*, 226:34–74, 2015.

[184] Markus Stumptner and Franz Wotawa. Coupling CSP decomposition methods and diagnosis algorithms for tree-structured systems. In Georg Gottlob and Toby Walsh, editors, *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9-15, 2003*, pages 388–393. Morgan Kaufmann, 2003.

[185] M. Tappler, A. Pferscher, B. K. Aichernig, and B. Könighofer. Learning and repair of deep reinforcement learning policies from fuzz-testing data. In *2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE)*, pages 27–39, Los Alamitos, CA, USA, apr 2024. IEEE Computer Society.

[186] Martin Tappler, Bernhard K. Aichernig, Giovanni Bacci, Maria Eichlseder, and Kim G. Larsen. L^*-based learning of markov decision processes (extended version). *Formal Aspects of Computing*, 33(4-5):575–615, 2021. Open Access.

[187] Martin Tappler, Bernhard K. Aichernig, and Roderick Bloem. Model-based testing IoT communication via active automata learning. In *2017 IEEE International Conference on Software Testing, Verification and Validation, ICST 2017, Tokyo, Japan, March 13-17, 2017*, pages 276–287. IEEE Computer Society, 2017.

[188] Martin Tappler, Bernhard K. Aichernig, Kim Guldstrand Larsen, and Florian Lorber. Time to learn - learning timed automata from tests. In Étienne André and Mariëlle Stoelinga, editors, *Formal Modeling and Analysis of Timed Systems - 17th International Conference, FORMATS 2019, Amsterdam, The Netherlands, August 27-29, 2019, Proceedings*, volume 11750 of *Lecture Notes in Computer Science*, pages 216–235. Springer, 2019.

[189] Martin Tappler, Filip Cano Córdoba, Bernhard K. Aichernig, and Bettina Könighofer. Search-based testing of reinforcement learning. In Luc De

Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 503–510. ijcai.org, 2022.

[190] Thi Ngoc Trang Tran, Alexander Felfernig, and Nava Tintarev. Humanized recommender systems: State-of-the-art and research issues. *ACM Trans. Interact. Intell. Syst.*, 11(2), jul 2021.

[191] Thi Ngoc Trang Tran, Alexander Felfernig, Christoph Trattner, and Andreas Holzinger. Recommender systems in the healthcare domain: state-of-the-art and research issues. *Journal of Intelligent Information Systems*, 57:171–201, 2018.

[192] F. Wallner, B. K. Aichernig, and C. Burghard. It's not a feature, it's a bug: Fault-tolerant model mining from noisy data. In *2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE)*, pages 316–328, Los Alamitos, CA, USA, apr 2024. IEEE Computer Society.

[193] Johannes P. Wallner, Andreas Niskanen, and Matti Järvisalo. Complexity results and algorithms for extension enforcement in abstract argumentation. *Journal of Artificial Intelligence Research*, 60:1–40, 2017.

[194] Jörg Weber and Franz Wotawa. Diagnosis and repair of dependent failures in the control system of a mobile autonomous robot. *Appl. Intell.*, 36(3):511–528, 2012.

[195] Martin Weiglhofer, Bernhard Aichernig, and Franz Wotawa. Fault-based conformance testing in practice. *International Journal of Software and Informatics*, 3(2–3):375–411, June/September 2009. Special double issue on Formal Methods of Program Development edited by Dines Bjoerner.

[196] Martin Weiglhofer, Bernhard K. Aichernig, and Franz Wotawa. Fault-based conformance testing in practice. *Int. J. Softw. Informatics*, 3(2-3):375–411, 2009.

[197] Martin Weiglhofer, Gordon Fraser, and Franz Wotawa. Using coverage to automate and improve test purpose based testing. *Inf. Softw. Technol.*, 51(11):1601–1617, 2009.

[198] T. White. *Hadoop: The Definitive Guide*. O'Reilly Media, Inc., 1st edition, 2009.

[199] Sandra Wilfling, Masoud Ebrahimi, Qamar Alfalouji, Gerald Schweiger, and Mina Basirat. Learning non-linear white-box predictors: A use case in energy systems. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 507–512. IEEE, 2022.

[200] Sandra Wilfling, Basak Falay, Qamar Alfalouji, and Gerald Schweiger. A dymola-python framework for data-driven model creation and co-simulation. In *Modelica Conferences*, pages 165–170, 2022.

[201] W. Eric Wong, Ruizhi Gao, Yihao Li, Rui Abreu, and Franz Wotawa. A survey on software fault localization. *IEEE Trans. Software Eng.*, 42(8):707–740, 2016.

[202] Franz Wotawa. Debugging hardware designs using a value-based model. *Appl. Intell.*, 16(1):71–92, 2002.

[203] Franz Wotawa. On the relationship between model-based debugging and program slicing. *Artif. Intell.*, 135(1-2):125–143, 2002.

[204] Franz Wotawa. Debugging VHDL designs: Introducing multiple models and first empirical results. *Appl. Intell.*, 21(2):159–172, 2004.

[205] Franz Wotawa. On the use of abduction as an alternative to decision trees in environmental decision support systems. *Int. J. Agric. Environ. Inf. Syst.*, 2(1):63–82, 2011.

[206] Franz Wotawa and David Kaufmann. Model-based reasoning using answer set programming. *Appl. Intell.*, 52(15):16993–17011, 2022.

[207] Franz Wotawa, Mihai Nica, and Iulia Moraru. Automated debugging based on a constraint model of the program and a test case. *J. Log. Algebraic Methods Program.*, 81(4):390–407, 2012.

[208] Franz Wotawa and Ingo Pill. On classification and modeling issues in distributed model-based diagnosis. *AI Commun.*, 26(1):133–143, 2013.

[209] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: Cluster computing with working sets. In Erich M. Nahum and Dongyan Xu, editors, *2nd USENIX Workshop on Hot Topics in Cloud Computing, HotCloud'10, Boston, MA, USA, June 22, 2010*. USENIX Association, 2010.

[210] Safdar Zaman, Gerald Steinbauer, Johannes Maurer, Peter Lepej, and Suzana Uran. An integrated diagnosis and repair architecture for ros-based robot systems. In *International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, 2013. IEEE.

[211] Xinhai Zhang, Jianbo Tao, Kaige Tan, Martin Törngren, José Manuel Gaspar Sánchez, Muhammad Rusyadi Ramli, Xin Tao, Magnus Gyllenhammar, Franz Wotawa, Naveen Mohan, Mihai Nica, and Hermann Felbinger. Finding critical scenarios for automated driving systems: A systematic mapping study. *IEEE Trans. Software Eng.*, 49(3):991–1026, 2023.

[212] Martin Zimmermann and Franz Wotawa. An adaptive system for autonomous driving. *Softw. Qual. J.*, 28(3):1189–1212, 2020.

[213] Martin Zimmermann, Franz Wotawa, and Ingo Pill. Pursuing intelligent behavior in cyber-physical systems by lightweight diagnosis. *Adv. Intell. Syst.*, 4(4), 2022.

WISSEN ▪ TECHNIK ▪ LEIDENSCHAFT

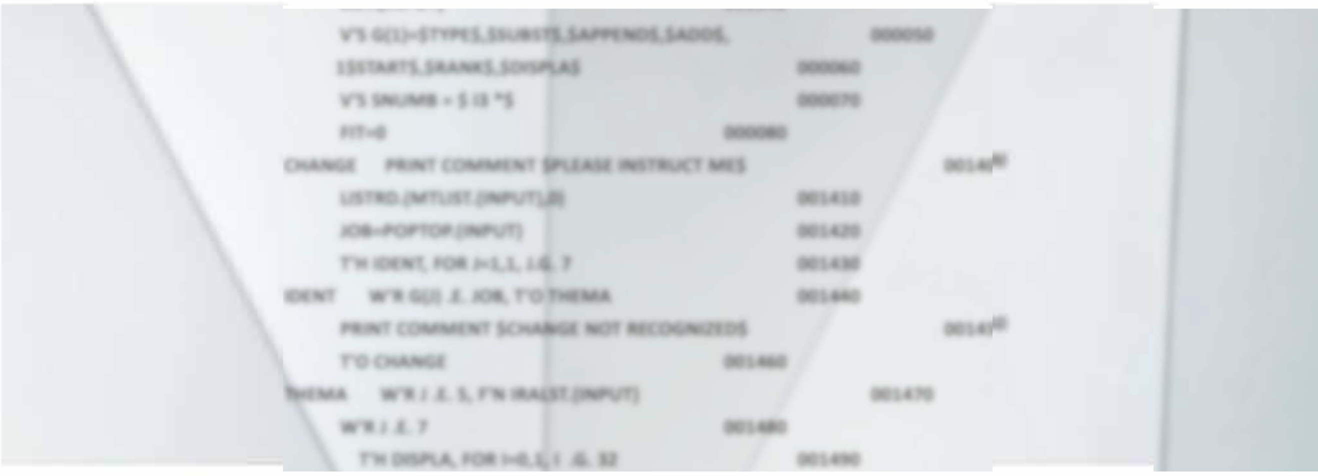

Technische Universität Graz
Institute of Software Engineering and Artificial Intelligence
Inffeldgasse 16b/2, 8010 Graz, Austria
T: +43(0)316/873-5711
office.sai@tugraz.at
▶ www.tugraz.at/institute/sai/home/

Verlag der Technischen Universität Graz
www.tugraz-verlag.at
ISBN (print) 978-3-85125-977-3
ISBN (e-book) 978-3-85125-978-0
DOI 10.3217/978-3-85125-977-3

The TU Graz Institute of Software Technology - The First 20 Years