Suhas Thejaswi

Curriculum Vitae

thejaswi@mpi-sws.org ◇ https://suhastheju.github.io
Nationality: India 🇮🇳 ◇ *Permanent Residence*: Finland 🇫🇮

Academic Positions

Apr 2023 – *present* — **Max Planck Institute for Software Systems** (Kaiserslautern, Germany 🇩🇪)
Postdoctoral Fellow

Jan 2023 – Oct 2024 — **Aalto University** (Espoo, Finland 🇫🇮)
Academic Visiting Fellow

Nov 2020 – Dec 2022 — **Aalto University** (Espoo, Finland 🇫🇮)
Postdoctoral Fellow

Industry Positions

Jan 2014 – Aug 2014 — **Nokia Networks** (Bangalore, India 🇮🇳)
Research Engineer

May 2012 – Jan 2014 — **Motorola Solutions** (Bangalore, India 🇮🇳)
Software Engineer

Mar 2010 – May 2012 — **Wipro Technologies** (Bangalore, India 🇮🇳)
Software Engineer

Education

2018 – 2022 — **Aalto University** (Espoo, Finland 🇫🇮)
PhD in Computer Science

2014 – 2017 — **Aalto University** (Espoo, Finland 🇫🇮)
MSc in Computer Science

2005 – 2009 — **Visveswaraya Technological University** (Belgaum, India 🇮🇳)
BEng in Computer Science

Research Experience

Apr 2023 – *present* — **Max Planck Institute for Software Systems**, *Postdoctoral Fellow*
Advisor: *Dr. Manuel Gomez-Rodriguez*
Research in Human-Machine collaboration. In assisted decision-making tasks, humans rely on algorithmic predictions to make final decisions, with the goal of achieving human-AI complimentary, where the efficiency of the human-AI partnership exceeds that of humans or algorithmic decisions independently. Designing such systems require understanding of human-behavior, as how humans use algorithmic predictions in their decision-making processes and addressing concerns like the emergence and propagation of biases from human decisions to human-AI collaborative decisions. My research introduces conceptual innovations that pave the way for a new paradigm of decision-making, where human expertise and algorithmic predictions are integrated with a critical awareness of the possibility to perpetuate human biases.

Research Experience (continued)

Oct 2021 – Dec 2022 — **Department of Computer Science, Aalto University**, *Postdoctoral Fellow*

Mar 2018 – Oct 2022 — **Department of Computer Science, Aalto University**, *Doctoral Researcher*
Thesis: *"Scalable Algorithm Designs for Mining Massive Datasets"*
Advisor: *Prof. Aristides Gionis*
My PhD research focused on the study of problems motivated from social issues, such as, modeling epidemic propagation using time-evolving graphs and mitigating bias in algorithmic decision-making systems. I pursued this goal by designing and engineering of algorithms with provable theoretical guarantees while emphasizing the empirical scalability to massive real-world datasets, and developing ethically responsible algorithmic techniques that ensure fairness in decision-making systems. My PhD thesis contributed to design of scalable algorithmic frameworks to solve important problems such as contact-tracing in epidemic models, avoiding bias in algorithmic decision making systems and providing theoretical insights to distinguish between fairness notions that are (not) possible to achieve in practice.

Apr 2015 – Nov 2018 — **Department of Computer Science, Aalto University**, *MSc Student*
Thesis: *"Scalable Parameterized Algorithms for Two Steiner Problems"*
Advisor: *Prof. Petteri Kaski*
Conducted research on design and engineering of parameterized algorithms for the Steiner tree and group Steiner tree problems—finding a smallest subgraph connecting a given set of vertices. Developed implementations optimized for parallel computing architectures, capable of scaling to graphs with billions of edges. This work secured sixth place in the Parameterized Algorithms and Exact Computation Experiments Challenge (PACE 2018).

Jul 2015 – Mar 2018 — **Department of Computer Science, Aalto University**, *Research Assistant*
Project: *"Algorithm Engineering for High-Performance Computing"*
Advisor: *Prof. Petteri Kaski*
Design of scalable algorithms for pattern detection in graphs, and presenting practical implementations that achieves near-peak arithmetic and memory bandwidth utilization on a range of vector parallel micro-architectures such as General Purpose Graphical Processing Units (GPGPUs). Also worked on micro-benchmarking of GPGPUs to identify performance bottlenecks. This work led to release of an *open source* software for motif search—identifying connected subgraph patterns in a graph—that can scale to large graphs as well as large patterns.

Aug 2009 – Feb 2010 — **Defense Research and Development Organization, India**, *Research Intern*
Project: *"Secure File-transfer Application with Multi-level Authentication"*
Advisor: *Sosale Guruprasad Gopinath*
Developed a secure file transfer application with a multi-factor authentication system for file transfers.

Industry Experience

Jan 2014 – Aug 2014 — **Nokia Networks**, *Research Engineer*

May 2012 – Jan 2014 — **Motorola Solutions**, *Software Engineer*

Industry Experience (continued)

Mar 2010 – May 2012 — **Ericsson Offshore Development Center**, *Software Engineer*
Contributed to the design and development of telecom call-processing systems at leading telecom equipment vendors such as Ericsson, Motorola and Nokia. Gained experience in telecom network management at Ericsson (as a consultant from Wipro Technologies), worked on Push-to-Talk services at Motorola, and Internet Multimedia Subsystems at Nokia.

Awards and Recognition

2025 — **Top Reviewer Award**
Recognized as a top reviewer at the International Conference on Machine Learning (ICML 2025).

Outstanding Reviewer Award
Recognized as an outstanding reviewer at the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD, 2025).

2022 — **European Union Transnational Access Scholarship**
Through a competitive selection process, awarded €5000 grant for research travel by the SoBigData++ consortium.

Aalto Foundation Grant
Received €2000 grant for covering PhD defense expenses.

2 × **HIIT Travel Grant**
Recieved €8000 grant for supporting research visit(s).

2021 — **Nokia Foundation Scholarship**
Through a competitive selection process, awarded €6000 by Nokia Foundation.

2020 — **SIAM Student Travel Grant**
Awarded $2000 travel grant for covering the conference expenses.

2019 — **TKK Scholarship**
Received €6000 travel grant for research visits.

2018 — **6th place in PACE programming contest**
Achieved 6th place in Parameterized Algorithms and Computational Experiments challenge.

2014 — **Moments Award**
Awarded by Motorola Corporation for excellent work in networks division.

2011 — **Connoisseur Award**
Awarded by Ericsson offshore development center for exemplary technical proficiency.

In the Media

Aug 2022 — **ACM Kudos research showcase** — Computational challenges of inclusivity and diversity in algorithmic decision making.

Teaching Experience

Fall 2023/24 — **Max Planck Institute for Software Systems**, *Instructor*
Module: *Seminar Course on Human Centric Machine Learning*
Responsibilities: Designing seminar content, assessed student performance, and office hours.

Teaching Experience (continued)

Spring 2021 — **Department of Computer Science, Aalto University**, *Teaching Assistant*
Module: *Spectral Graph Theory and Signed Graphs*
Responsibilities: Assessed student performance, assisted with creating course content and office hours.

Module: *Programming Parallel Computers* – **Best course award Runner-up**
Responsibilities: Graded assignments, conducted lab sessions to support student learning.

Fall 2018/19 — **Department of Computer Science, Aalto University**, *Teaching Assistant*
Module: *Algorithmic Methods of Data Mining*
Responsibilities: Delivered lectures, evaluated student performance through assignments and exams, and held office hours. Also responsible for creating assignments and projects that complemented lecture material.

Spring 2018 — **Department of Computer Science, Aalto University**, *Teaching Assistant*
Module: *Modern Database Systems*
Responsibilities: Evaluated student performance through assignments and exams, and held office hours.

Fall 2016 — **Department of Computer Science, Aalto University**, *Teaching Assistant*
Module: *Principles of Algorithmic Techniques*
Responsibilities: Assisted in designing assignments and exams, evaluated student performance, and conducted help sessions to support student learning.

Research Publications

All publications are presented in **reverse chronological order**.

Disclaimer: In Computer Science, particularly in the areas of Machine Learning and Data Mining, high-quality conference publications are often more highly regarded than journal articles. Additionally, when the articles are primarily theoretical, it is customary to list authors in alphabetical order, and these articles are indicated by $\alpha\beta$.

Note on the provided rankings: Conference rankings are reported according to the CORE rankings as of the date of publication for each entry. These rankings categorize conferences as follows:

- **A*** – flagship conference and a leading venue within a discipline,
- **A** – excellent conference and highly respected within a discipline,
- **B** – good to very good conference and well-regarded within a discipline,
- **C** – other recognized conferences that meet minimum standards.

Table I: *Summary of conference publications over time*

Time frame	Conference ranks
Postdoctoral research	A* A* A* A* A*
Doctoral research	A* A A B **Best paper candidate**

Manuscripts under review

1. Authors: Nina Corvelo-Benz, Stratis Tsirtsis, Eleni Straitouri, Ivi Chatzi, Ander Artola Velasco, **Suhas Thejaswi**, and Manuel Gomez-Rodriguez
   Title: *"Evaluation of Large Language Models via Coupled Token Generation"*
   Preprint: https://arxiv.org/abs/2502.01754
   Year: 2025

2. Authors:[αβ] Ameet Gadekar and **Suhas Thejaswi**
   Title: *"Capacitated fair-range clustering: hardness and approximation algorithms"*
   Preprint: https://arxiv.org/abs/2401.05502
   Year: 2025

3. Authors: **Suhas Thejaswi**, Ameet Gadekar, Bruno Ordozgoiti, and Aristides Gionis
   Title: *"Diversity-aware clustering: computational complexity and approximation algorithms"*
   Preprint: https://arxiv.org/abs/2401.05502
   Year: 2024

Peer-reviewed journal articles

1. Title: *"Restless Reachability Problems in Temporal Graphs"*
   Authors: **Suhas Thejaswi**, Juho Lauri, and Aristides Gionis
   Journal: Knowledge and Information Systems
   Year: 2025
   Volume: /
   Impact Fac.: 2.6

2. Title: *"Finding Path Motifs in Large Temporal Graphs Using Algebraic Fingerprints"*
   Authors: **Suhas Thejaswi**, Aristides Gionis, and Juho Lauri
   Journal: Big Data
   Year: 2020
   Volume: 8 / 5
   Impact Fac.: 12.4
   Notes: *Special issue on best papers of SIAM Data Mining* 2020

Peer-reviewed conference proceedings

1. Title: *"Fair clustering for data summarization: improved approximation algorithms and complexity insights"*
   Authors:[αβ] Ameet Gadekar, Aristides Gionis, and **Suhas Thejaswi**
   Venue: The ACM Web Conference: Research Track (WWW)
   Year: 2025
   CORE rank: A*
   Accept. rate: 20%

2. Title: *"Matchings, Predictions and Counterfactual Harm in Refugee Resettlement Processes"*
   Authors: Seungeon Lee, Nina Corvelo-Benz, **Suhas Thejaswi**, and Manuel Gomez-Rodriguez
   Venue: Causal Learning and Reasoning (CLeaR)
   Year: 2025
   CORE rank: Not available
   Accept. rate: 43.94%

3. Title: *"Prediction-Powered Ranking of Large Language Models"*
   Authors: Ivi Chatzi, Eleni Straitouri, **Suhas Thejaswi**, and Manuel Gomez Rodriguez
   Venue: Advances in Neural Information Processing Systems (NeurIPS)
   Year: 2024
   CORE rank: A*
   Accept. rate: 25.8%

4 Title: *"Fair Column Subset Selection"*
Authors: Antonis Matakos, Bruno Ordozgoiti, and **Suhas Thejaswi**
Venue: ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)
Year: 2024
CORE rank: A*
Accept. rate: 20%

5 Title: *"Controlling Counterfactual Harm in Decision Support Systems Based on Prediction Sets"*
Authors: Eleni Straitouri, **Suhas Thejaswi**, and Manuel Gomez Rodriguez
Venue: Advances in Neural Information Processing Systems (NeurIPS)
Year: 2024
CORE rank: A*
Accept. rate: 25.8%

6 Title: *"Towards Human-AI Complementarity with Predictions Sets"*
Authors: Giovanni De Toni, Nastaran Okati, **Suhas Thejaswi**, Eleni Straitouri, and Manuel Gomez-Rodriguez
Venue: Advances in Neural Information Processing Systems (NeurIPS)
Year: 2024
CORE rank: A*
Accept. rate: 25.8%

7 Title: *"Clustering with Fair-Center Representation: Parameterized Approximation Algorithms and Heuristics"*
Authors: **Suhas Thejaswi**, Ameet Gadekar, Bruno Ordozgoiti, and Michal Osadnik
Venue: ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)
Year: 2022
CORE rank: A*
Accept. rate: 15%
Notes: **Featured in ACM research showcase**

8 Title: *"Diversity-Aware k-median: Clustering with Fair Center Representation"*
Authors: **Suhas Thejaswi**, Bruno Ordozgoiti, and Aristides Gionis
Venue: Machine Learning and Knowledge Discovery in Databases (ECML)
Year: 2021
CORE rank: A
Accept. rate: 19.1%

9 Title: *"Pattern detection in large temporal graphs using algebraic fingerprints"*
Authors: **Suhas Thejaswi** and Aristides Gionis
Venue: SIAM International Conference on Data Mining (SDM)
Year: 2020
CORE rank: A
Accept. rate: 24%
Notes: **Best paper candidate**

10 Title: *"Engineering Motif Search for Large Motifs"*
Authors:[αβ] Petteri Kaski, Juho Lauri, and **Suhas Thejaswi**
Venue: International Symposium on Experimental Algorithms (SEA)
Year: 2018
CORE rank: B
Accept. rate: Not available

Peer-reviewed workshop papers

1
Title: *"Evaluation of Large Language Models via Coupled Token Generation"*
Authors: Nina Corvelo-Benz, Stratis Tsirtsis, Eleni Straitouri, Ivi Chatzi, Ander Artola Velasco, **Suhas Thejaswi**, and Manuel Gomez-Rodriguez
Venue: Workshop on Building Trust in LLMs and LLM Applications
Year: 2025
Notes: *Building Trust @ International Conference on Learning Representation (ICLR)*

2
Title: *"Prediction Powered Ranking of Large Language Models"*
Authors: Ivi Chatzi, Eleni Straitouri, **Suhas Thejaswi**, and Manuel Gomez-Rodriguez
Venue: Workshop on Human-centered Evaluation and Auditing of Language Models (HEAL @ CHI)
Year: 2024
Notes: *HEAL @ ACM Conference on Human Factors in Computing Systems (HCI)*

3
Title: *"Matchings, Predictions and Counterfactual Harm in Refugee Resettlement Processes"*
Authors: Seungeon Lee, Nina Corvelo-Benz, **Suhas Thejaswi**, and Manuel Gomez-Rodriguez
Venue: Workshop on Ethical Artificial Intelligence: Methods and Applications (EAI)
Year: 2024
Notes: *EAI @ ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*

4
Title: *"Controlling Counterfactual Harm in Decision Support Systems Based on Prediction Sets"*
Authors: Eleni Straitouri, **Suhas Thejaswi**, and Manuel Gomez Rodriguez
Venue: Workshop on Humans, Algorithmic Decision-Making and Society (HADS)
Year: 2024
Notes: *HADS @ International Conference on Machine Learning (ICML)*

Theses

1
Title: *"Scalable Algorithm Designs for Mining Massive Datasets"*
Author: **Suhas Thejaswi**
School: Aalto University
Year: 2022
Type: **PhD thesis**
Series: Aalto University Press
ISBN: ISBN: 978-952-64-0942-9

2
Title: *"Scalable Parameterized Algorithms for Two Steiner Problems"*
Author: **Suhas Thejaswi**
School: Aalto University
Year: 2018
Type: **MSc thesis**
Series: Aalto University Press

Academic Service

Conferences and journals

2025 **Program Committee Member**
– *ACM Conference on Knowledge Discovery and Data Mining (KDD)*
– *ACM Conference on Web Search and Knowledge Discovery (WSDM)*

Academic Service (continued)

Reviewer
– *Frontiers of Computer Science*
– *Advances in Neural Information Processing Systems (NeurIPS)*
– *International Conference on Machine Learning (ICML)*
– *ACM International World Wide Web Conference (WebConf)*
– *Journal of Information and Computation*

2024 **Program Committee Member**
– *ACM Conference on Knowledge Discovery and Data Mining (KDD)*
– *ACM Conference on Web Search and Knowledge Discovery (WSDM)*
– *International Conference on Artificial Intelligence and Statistics (AISTATS)*

Reviewer
– *Advances in Neural Information Processing Systems (NeuRIPS)*
– *Journal of Artificial Intelligence (JIA)*
– *Journal of Machine Learning Research (JMLR)*

2023 **Program Committee Member**
– *ACM Conference on Knowledge Discovery and Data Mining (KDD)*
– *ACM Conference on Web Search and Knowledge Discovery (WSDM)*

Reviewer
– *Symposium on Theoretical Aspects of Computer Science (STACS)*
– *ACM Transactions of Knowledge Discovery from Data (TKDD)*
– *ACM Transactions on the Web (TWEB)*

2022 **Reviewer**
– *European Symposium of Algorithms (ESA)*
– *ACM Transactions of Knowledge Discovery from Data (TKDD)*

2021 **Reviewer**
– *Algorithmica Journal*
– *ACM Transactions of Knowledge Discovery from Data (TKDD)*

2019 **Reviewer**
– *Symposium on Experimental Algorithms (SEA)*

2018 **Local Organizing Committee Member**
– *European Symposium of Algorithms (ESA)*

2007 – 2009 **Student-chair of Editorial Board – "Bits-n-Bytes"**
Initiated the student edition of the technical newsletter "Bits-n-Bytes" during undergraduate studies, overseeing the publication process and responsible for launching the first edition of the newsletter. Also, chaired the editorial board from 2007 to 2009.

University administration

2018 – 2019 **University Student Admissions @ Aalto University**
Assisted in shortlisting student applicants for MSc in computer science program at Aalto University for the academic years 2018 and 2019.

Public engagement and outreach

2024 – *present* **Mentorship for Relief @ Max Planck Institute for Software Systems**
Assisted in creation of a working group on an initiative to support researchers in areas affected by war and conflicts.

Academic Service (continued)

Jul 2023 — **Panel member in discussion on research in academia Vs industry**
Part of the panel discussion on research in academia Vs research in industry at the International Max Planck Research Summer School, 2023 held at Saarbrucken, Germany.

Talks and Presentations

Jul 2024 & Aug 2024 — **Matchings, Predictions and Counterfactual Harm in Refugee Resettlement Processes**
– Talk @ KTH Royal Institute of Technology, Sweden
– Workshop on Ethical Aspects of Artificial Intelligence co-located at KDD (EAI-KDD), Barcelona, Spain.

Nov 2023 — **Clustering with Fair Center Representation: Theory and Practice**
– Talk @ Max Planck Institute for Software Systems, Kaiserslautern, Germany.

Jul 2022 — **Clustering with Fair Center Representation: Parameterized Approximation Algorithms and Heuristics**
– ACM Conference on Data Mining (KDD 2022), Washington DC, USA.

Oct 2021 & Jul 2022 — **Finding Path Motifs in Temporal Graphs using Algebraic Fingerprints**
– Talk @ Technical University of Berlin.
– Workshop on Algorithmic Aspects of Temporal Graphs co-located with ICALP (AATG-ICALP, 2022), Paris, France.

Jul 2021 — **Parameterized Algorithms: from Theory to Practice**
– Talk @ Aalto University, Finland.

Sep 2021 & Dec 2021 — **Diversity-aware k-median: Clustering with fair center representation**
– Artificial Intelligence Day, Aalto University (AI day 2021), Finland
– European Conference on Machine Learning (ECML 2021), Bilbao, Spain.

May 2020 & Nov 2019 — **Pattern Detection in Large Temporal Graphs Using Algebraic Fingerprints**
– SIAM Conference on Data Mining (SDM 2020), Cincinatti, USA
– Artificial Intelligence Day, Aalto University, Finland.

Aug 2018 — **Engineering Motif Search for Large Motifs**
– Symposium of Experimental Algorithms (SEA 2018).

Miscellaneous Experience

Academic Visits

Jun 2024 – Aug 2024 — **KTH Royal Institute of Technology** (**KTH**), Stockholm, Sweden
Visited Prof. Aristides Gionis and the data-mining group at KTH.

May 2022 – Jun 2022 — **Queen Mary University of London** (**QMUL**), London, United Kingdom
Visited Dr. Bruno Ordozgoiti and his research group at QMUL.

Oct 2021 – Oct, 2021 — **Technical University of Berlin** (**TU-Berlin**), Berlin Germany.
Visited Professor Rolf Niedermier, and the algorithms and complexity research group at TU Berlin.

Miscellaneous Experience (continued)

Summer Schools

Aug 2023 — **Max Planck Advanced Course on Foundations of Computer Science** (**ADFOCS**-2023), Saarbrucken, Germany.
Topic(s): Algorithmic Foundations of Data Analysis; Clustering

Oct 2021 — **Advanced course on Human Centered AI**, Berlin, Germany
Topic(s): Societal, legal and ethical impact of AI; Learning with human in the loop; Human computer interaction

Jul 2019 — **Gdansk Summer School on Parameterized algorithms**, Gdansk, Poland.
Topic(s): Parameterized algorithms; Approximation algorithms

Aug 2019 — **Max Planck Advanced Cource on Foundations of Computer Science** (**ADFOCS** 2019), Saarbrucken, Germany.
Topic(s): Distributed algorithms; Game theory fundamentals

Jul 2015 — **Summer school on High Performance Computing**, Espoo, Finland.
Topic(s): Message passing interface (MPI); OpenMP; CUDA programming

Open-source Software

Pattern detection — **Temporal patterns:** An algebraic algorithm framework for finding path motifs in time-evolving graphs.
– https://github.com/suhastheju/temporal-patterns
– https://github.com/suhastheju/temporal-patterns-mk2

Restless reachability: An algebraic algorithm framework for finding restless paths in time-evolving graphs.
– https://github.com/suhastheju/restless-reachability

Graph motifs: A CUDA implementation of an algebraic framework for finding connected subgraphs in static graphs.
– https://github.com/suhastheju/motif-localized

Steiner tree — **Erickson-Monma-Veinott algorithm:** A thread-parallel implementation using OpenMP and bit-twiddling hacks to enable fast subset enumeration, capable of scaling graphs with billion edges.
– https://github.com/suhastheju/pace-2018-exact

Dreyfus-Wagner Algorithm: A thread-parallel implementation using OpenMP.
– https://github.com/suhastheju/steiner-dreyfus-wagner

Shortest path — **Dijkstra's algorithm:** A highly scalable single-threaded implementation of the single-source shortest path algorithm for edge-weighted graphs, capable of handling graphs with billions of edges.
– https://github.com/suhastheju/shortest-path

Clustering — **Fair-clustering:** Implementation of a collection of algorithms for fair clustering problems with k-median, k-means, k-center and k-supplier objectives.
– https://github.com/suhastheju/diversity-aware-clustering

Skills

Programming — Proficient – C, C++, Python
Course work – Matlab, R

Parallel Computing — Proficient – CUDA, OPENMP, MPI
Course work – MapReduce, Hadoop, Spark

Skills (continued)

ML Libraries	Proficient – Pandas, Numpy, Scipy, Matplotlib Course work – Scikit-Learn, TensorFlow
Operating Systems	Proficient – Unix, Linux
Databases	Course work – PSQL, MongoDB
Misc.	Make system, LaTeX, CSP, SAT, ILP, MIP Solvers