Audris Mockus

Department of Electrical Engineering and Computer Science
Min H. Kao Building, Rm 613
1520 Middle Drive
Knoxville, TN, 37996-2250

Phone: +1 865 974 2265
Fax: +1 865 974 5483
audris@utk.edu, http://mockus.org
Google Scholar: h-i: 57, i10-i: 126

Personal

Born July 14, 1964; in Kaunas, Lithuania
Lithuanian Citizen; US permanent resident

Education

Ph.D. (Statistics), 1994, Carnegie Mellon University
M.S. (Statistics), 1991, Carnegie Mellon University
B.S. & M.S. (Automatic Control), 1988,
Moscow Institute of Physics and Technology

Positions

Department of Electrical Engineering and Computer Science, University of Tennessee,
Knoxville, TN 2014- the Harlan D. Mills Chair Professor

on sabbatical visiting META Inc,
New York, NY 2022- Research Scientist

on sabbatical visiting Vilnius University,
Vilnius, Lithuania 2022-2023 Research Scientist

Avaya Labs Research,
Basking Ridge, NJ 2001-2016 Consulting Research Scientist

Bell Laboratories, Software Production Research Department,
Naperville, IL 1995-2001 Member of Technical Staff

Carnegie Mellon University, Dept. of Statistics,
Pittsburgh, PA 1994-1995 Visiting Assistant Professor

Carnegie Mellon University, Dept. of Statistics,
Pittsburgh, PA 1990-1994 Teaching and Research Assistant

Institute of Mathematics and Informatics,
Vilnius, Lithuania 1988-1990 Research Fellow

Kaunas University of Technology,
Kaunas, Lithuania 1989-1990 Teaching Assistant

Computer Center of the Academy of Sciences of the U.S.S.R.,
Moscow, Russia 1985-1988 Research Assistant

Recognition

Foundational Contribution Award, Mining Software Repositories (MSR) 2023.

Honorable Mention, Most Influential Paper Award of International Conference on Software Engineering 2001. Awarded in 2011.

Most influential paper of International Conference on Software Engineering 2000. Selected by a specially constituted committee of ICSE 2010 using a range of indicators of influence, impact, and significance.

International Conference on Software Engineering, 2025, best paper award in the industry track.

Large Language Models for Code (LLM4C), 2025, best paper award.

International Conference on Software Engineering 2025, distinguished paper award in the main track.

International Conference on Software Engineering, 2023, distinguished reviewer.

ACM SIGSOFT International Symposium on the Foundations of Software Engineering, 2018, distinguished paper award.

Mining Software Repositories, 2014, distinguished paper award.

Foundations of Software Engineering, 2010, distinguished paper award.

Ranked ninth on the list of the top software engineering scholars in the peer reviewed article in Communications of the Association of Computing Machinery, 2007.

First International Workshop on Emerging Trends in FLOSS Research and Development, 2007, best paper award.

International Conference on Software Engineering, 2000, best paper award.

International Conference on Software Maintenance, 2000, best paper award.

Grants and Patents

1. PI: Audris Mockus, Co-PI: Dawnie Steadman Sponsor: National Institute Of Justice, Deep Learning Methods for Post Mortem Interval Estimation, 15PNIJ-21-GG-04161-SLFO, Jul 2022 - Dec 2024 Amount: 406K
2. PI: Audris Mockus, Co-PI: Jian Huang, James Herbsleb Sponsor: National Science Foundation Collaborative Research: CCRI: New: World Of Code (WoC): The development of curated code resource to support research in software engineering, CRI-2120429, Oct 2021 - Sep 2023 Amount: 712K
3. PI: Audris Mockus, Co-PI: Jian Huang, James Herbsleb Sponsor: National Science Foundation CCRI: Planning: Collaborative Research: World Of Code (WoC): Infrastructure for Open Source Census, CRI-1925615, Oct 2019 - Sep 2020 Amount: 100K
4. co-PI: Audris Mockus, PI: James Herbsleb, Co-PI: Bogdan Vasilescu Sponsor: National Science Foundation CHS: Medium: Collaborative Research: SDI-CPR: Sustaining Digital Infrastructure as a Common Pool Resource. IIS-1901102, Oct 2019 - Sep 2023 Amount: 399K Total: 800K
5. PI: Audris Mockus, Co-PI: Dawnie Steadman, Automation-Supported Curation of Large Forensic Image Databases Sponsor: National Institute Of Justice, USDOJ-2018-DU-BX-0181, Jan 2018 - Dec 2020 Amount: 354,751
6. PI: Audris Mockus, Co-PIs: Bogdan Bichecu, Randy Bradley, James Herbsleb, Russell Zaretzky Sponsor: National Science Foundation BIGDATA: Collaborative Research: IA: OSCAR - Open Source Supply Chains and Avoidance of Risk: An evidence based approach to improve FLOSS supply chains, IIS-1633437, Sep 2016 - Aug 2020 Amount: 665,097.00, total: 1,300,000
7. PI: Audris Mockus, Co-PI: Dawnie Steadman ICPUTRD: Image Cloud Platform for Use in Tagging and Re- search on Decomposition Sponsor: National Institute Of Justice, USDOJ-NIJ-2016-DN-BX-0179, Jan 2016 - Dec 2018 Amount: 327,372.00
8. PI: Audris Mockus A hybrid data structure for future ALM development Sponsor: UT-Battelle, LLC, UT-B 400015120, Oct 2016 - May 2017 Amount: 29,499.00
9. Audris Mockus Co-Principal Investigator Coordination, communication, and collaboration in open source software development NSF Grant 0414698, 03/01/06 – 05/31/09
10. Audris Mockus, Randy Hackbarth, John Palframan, and Brett Shockley. System and method for prioritizing customers and predicting service escalation. U.S. Patent 9176729, granted October 27 2015.
11. Audris Mockus, Randy Hackbarth, and John Palframan. System and method for prioritizing and remediating defect risk in source code. U.S. Patent 9172809, granted October 27 2015.
12. Kenneth Cox, Stacie Hibino, Lichan Hong, Audris Mockus, and Graham Wills. A method for graphically displaying an overview. U.S. Patent 8386920, 2013.
13. System and method for interactive visualization, analysis and control of a dynamic database. US Patent 5937064, granted on August 10, 1999

14. Method and apparatus for data visualization US Patent 8386920, granted on February 26, 2013

PhD Students

Sadika Amreen, Graduated on Apr, 2019, "Methods of Disambiguating andDe-anonymizing Authorship in LargeScale Operational Data"

Yuxing Ma, Graduated July, 2020, "Software Supply Chain (SSC) Development and Application"

Tapajit Dey, Graduated Dec, 2020, "Modeling User-affected Software Properties for Open Source Software Supply Chains"

Sara Mousavi, Graduated Fall, 2021, "Auto-curation of Large Evolving Image Datasets"

David Reid, Graduated Fall, 2023, "Applying the Universal Version History Concept to Help De-Risk Copy-Based Code Reuse"

Anna Maria Nau, Fall, 2024, "Applications of Machine Learning in Forensic Investigations of Human Decomposition"

Mahmoud Jahanshahi, Spring, 2025, "Copy-Based Reuse and its Implications in Open Source Software Supply Chains"

Patricia Jean Goedecke, Spring, 2025, "Social determinants of health with hypertensive disorders of pregnancy in Memphis, Tennessee: Retrospective cohort studies"

Articles in Refereed Proceedings

1. Audris Mockus, Peter C Rigby, Rui Abreu, Anatoly Akkerman, Yogesh Bhootada, Payal Bhuptani, Gurnit Ghardhora, Lan Hoang Dao, Chris Hawley, Renzhi He, Sagar Krishnamoorthy, Sergei Krauze, Jianmin Li, Anton Lunov, Dragos Martac, Francois Morin, Neil Mitchell, Venus Montes, Maher Saba, Matt Steiner, Andrea Valori, Shanchao Wang, and Nachiappan Nagappan. Metrics driven reengineering and continuous code improvement at meta. 2025.

2. Addi Malviya-Thakur, Reed Milewicz, Lavinia Paganini, Mahmoud Jahanshahi, Ahmed Samir Imam Mahmoud, Bogdan Vasilescu, and Audris Mockus. Scientific open-source software is more sustainable than one might think! In *The ACM International Conference on the Foundations of Software Engineering*, June 23-27 2025.

3. Mahmoud Jahanshahi, David Reid, Adam McDaniel, and Audris Mockus. Oss license identification at scale: A comprehensive dataset using world of code. In *2025 IEEE/ACM 22nd International Conference on Mining Software Repositories (MSR)*.

4. Rui Abreu, Vijayaraghavan Murali, Peter Rigby, Chandra Maddila, Weiyan Sun, Jun Ge, Kaavya Chinniah, Audris Mockus, Megh Mehta, and Nachi Nagappan. Moving faster and reducing risk: Using llms in release deployment. In *47th International Conference on Software Engineering*, April-May 2025. Best paper award.

5. Audris Mockus and Mahmoud Jahanshahi. Cracks in the stack: Hidden vulnerabilities and licensing risks in llm pre-training datasets. In *LLM4Code*, April-May 2025. Best paper award.

6. Courtney Miller, Mahmoud Jahanshahi, Audris Mockus, Bogdan Vasilescu, and Christian K astner. Understanding the response to open-source dependency abandonment in the npm ecosystem. In *47th International Conference on Software Engineering*, April-May 2025. Best paper award.

7. Mahmoud Jahanshahi, David Reid, Adam McDaniel, and Audris Mockus. Oss license identification at scale: A comprehensive dataset using world of code, 2024.

8. Mahmoud Jahanshahi and Audris Mockus. Dataset: Copy-based reuse in open source software. In *2024 IEEE/ACM 21st International Conference on Mining Software Repositories (MSR)*, pages 42–47. IEEE, 2024.

9. Addi Malviya-Thakur, Reed Milewicz, Lavinia Paganini, Ahmed Samir Imam Mahmoud, and Audris Mockus. Scicat: A curated dataset of scientific software repositories, 2023.

10. Anna-Maria Nau, Audris Mockus, and Dawnie Wolfe Steadman. Stage of decay estimation exploiting exogenous and endogenous image attributes to minimize manual labeling efforts and maximize classification performance. In *IEEE International Conference on Image Processing (ICIP)*, October 2023.

11. Audris Mockus, Peter C. Rigby, Rui Abreu, Parth Suresh, Yifen Chen, and Nachiappan Nagappan. Modeling the centrality of developer output with software supply chains. In *ESEC/FSE 2023*, December 2023.

12. David Reid and Audris Mockus. Applying the universal version history concept to help de-risk copy-based code reuse. In *23rd IEEE International Working Conference on Source Code Analysis and Manipulation (SCAM)*, October 2023.

13. addi Malviya-Thakur, Audris Mockus, Russel Zaretzki, Bogdan Bichescu, and Randy Bradley. How r developers explain their package choice: A survey. In *International Conference on Empirical Software Engineering*, October 2023.

14. Audris Mockus. Msr foundational contribution award. In *International Conference on Mining Software Repositories*, May 2023.

15. Audris Mockus, Alexander Nolte, and James Herbsleb. Msr mining challenge: World of code. In *Proceedings of the International Conference on Mining Software Repositories (MSR 2023)*, 2023.

16. David Reid, Mahmoud Janshani, and Audris Mockus. The extent of orphan vulnerabilities from code reuse in open source softwar. In *ICSE 2022*. ACM Press, May 2022.

17. Ahmed Samir Imam Mahmoud, Tapajit Dey, Alexander Nolte, Audris Mockus, and James D Herbsleb. One-off events? an empirical study of hackathon code creation and reuse. *Empirical Software Engineering*, 2022.

18. Ahmed Imam, Tapajit Dey, Alexander Nolte, Audris Mockus, and James D Herbsleb. The secret life of hackathon code where does it come from and where does it go? In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*, pages 68–79. IEEE, 2021.

19. Tapajit Dey, Andrey Karnauch, and Audris Mockus. Representation of developer expertise in opensource software. In *ICSE 2021*. ACM Press, May 2021.

20. Sara Mousavi, Dylan Lee, Tatianna Griffin, Kelley Cross, Dawnie Steadman, and Audris Mockus. Schism: Semantic clustering via image sequence merging for images of human-decomposition. In *WACV'21 Proceedings*, January 2021.

21. Sara Mousavi, Dylan Lee, Tatianna Griffin, Dawnie Steadman, and Audris Mockus. Collaborative learning of semi-supervised clustering and classification for labeling uncurated data. In *IEEE International Conference on Image Processing (ICIP)*, October 2020.

22. Andrey Krutauz, Tapajit Dey, Peter Rigby, and Audris Mockus. Do code review measures explain the incidence of post-release defects? case study replications and bayesian networks. In *ESEC/FSE'20 Journal First Track*, November 2020.

23. Tapajit Dey and Audris Mockus. Effect of technical and social factors on pull request quality for the npm ecosystem. In *Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement*, October 2020.

24. Audris Mockus, Diomidis Spinellis, Zoe Kotti, and Gabriel John Dusing. A complete set of related git repositories identified via community detection approaches based on shared commits. In *IEEE Working Conference on Mining Software Repositories: Data Showcase*, May 2020.

25. Tanner Fry, Tapajit Dey, Andrey Karnauch, and Audris Mockus. A dataset and an approach for identity resolution of 38 million author ids extracted from 2b git commits. In *IEEE Working Conference on Mining Software Repositories: Data Showcase*, May 2020.

26. Tapajit Dey, Bogdan Vasilescu, and Audris Mockus. An exploratory study of bot commits. In *Proceedings of the IEEE/ACM 42nd international conference on software engineering workshops*, pages 61–65, 2020.

27. Tapajit Dey, Sara Mousavi, Eduardo Ponce, Tanner Fry, Bogdan Vasilescu, Anna Filippova, and Audris Mockus. Detecting and characterizing bots that commit code. In *IEEE Working Conference on Mining Software Repositories*, May 2020.

28. Sara Mousavi, Dylan Lee, Tatianna Griffin, Dawnie Steadman, and Audris Mockus. An analytical workflow for clustering forensic images. In *AAAI CONFERENCE ON ARTIFICIAL INTELLIGENCE*, 2020.

29. Andrey Karnauch, Sadika Amreen, and Audris Mockus. Developer reputation estimator (dre). In *ASE'19*, 2019.

30. Sara Mousavi, Ramin Nabati, Megan Kleeschulte, Dawnie Steadman, and Audris Mockus. Machine-assisted annotation of forensic imagery. In *26th IEEE International Conference on Image Processing (ICIP)*, September 22-25 2019.

31. Tapajit Dey, Yuxing Ma, and Audris Mockus. Patterns of effort contribution and demand and user classification based on participation patterns in npm ecosystem. In *Proceedings of the 15th International Conference on Predictive Models and Data Analytics in Software Engineering*. ACM, 2019. Best paper award.

32. Yuxing Ma, Chris Bogart, Sadika Amreen, Russell Zaretzki, and Audris Mockus. World of code: An infrastructure for mining the universe of open source vcs data. In *IEEE Working Conference on Mining Software Repositories*, May 26 2019.

33. Sara Mousavi, Angela Dautart, Audris Mockus, and Dawnie W. Steadman. Machine learning to detect and localize forensics-relevant features. In *American Academy of Forensic Sciences, 2019 Annual Scientific Meetings*, page 445. AAFS, 2019.

34. Sara Mousavi, Angela Dautart, Audris Mockus, and Dawnie W. Steadman. Icputrd: Image cloud platform for use in tagging and research on decomposition. In *National Institute of Justice, 2019 Research Symposium*. NIJ, 2019.

35. Tapajit Dey and Audris Mockus. Are software dependency supply chain metrics useful in predicting change of popularity of npm packages? In *Proceedings of the 14th International Conference on Predictive Models and Data Analytics in Software Engineering*, pages 66–69. ACM, 2018.

36. Tapajit Dey and Audris Mockus. Modeling relationship between post-release faults and usage in mobile software. In *Proceedings of the 14th International Conference on Predictive Models and Data Analytics in Software Engineering*, pages 56–65, 2018.

37. Yuxing Ma, Tapajit Dey, and Audris Mockus. Modularizing global variable in climate simulation software: position paper. In *Proceedings of the International Workshop on Software Engineering for Science*, pages 8–11. ACM, 2016.

38. Wangcheng Yan, Audris Mockus, Tiffany B. Saul, and Dawnie W. Steadman. Curating forensic image collection using machine learning. In *American Academy of Forensic Sciences , 2018 ANNUAL SCIENTIFIC MEETING*, page 139. AAFS, 2018.

39. Minghui Zhou, Qingying Chen, Audris Mockus, and Fengguang Wu. On the scalability of linux kernel maintainers' work. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, pages 27–37. ACM, 2017.

40. Walton Macey, Dali Wang, Peter Thornton, and Audris Mockus. Live restructuring of data architecture: Wip. In *Proceedings of the 12th International Workshop on Software Engineering for Science*, pages 24–25. IEEE Press, 2017.

41. Nathan Wilder, Jared M Smith, and Audris Mockus. Exploring a framework for identity and attribute linking across heterogeneous data systems. In *Proceedings of the 2nd International Workshop on BIG Data Software Engineering*, pages 19–25. ACM, 2016.

42. Patrick Tendick and Audris Mockus. Decisions as a service for application centric real time analytics. In *2016 IEEE/ACM 2nd International Workshop on Big Data Software Engineering (BIGDSE)*, pages 1–7. IEEE, 2016.

43. Tapajit Dey, Jacob Logan Massengill, and Audris Mockus. Analysis of popularity of game mods: A case study. In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, pages 133–139. ACM, 2016.

44. Peter C Rigby, Yue Cai Zhu, Samuel M Donadelli, and Audris Mockus. Quantifying and mitigating turnover-induced knowledge loss: case studies of chrome and a project at avaya. In *Software Engineering (ICSE), 2016 IEEE/ACM 38th International Conference on*, pages 1006–1016. IEEE, 2016.

45. Jiaxin Zhu, Minghui Zhou, and Audris Mockus. Effectiveness of code contribution: From patch-based to pull-request-based tools. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 871–882. ACM, 2016.

46. Kazuhiro Yamashita, Changyun Huang, Meiyappan Nagappan, Yasutaka Kamei, Audris Mockus, Ahmed E Hassan, and Naoyasu Ubayashi. Thresholds for size and complexity metrics: A case study from the perspective of defect density. In *Software Quality, Reliability and Security (QRS), 2016 IEEE International Conference on*, pages 191–201. IEEE, 2016.

47. Meiyappan Nagappan, Romain Robbes, Yasutaka Kamei, Éric Tanter, Shane McIntosh, Audris Mockus, and Ahmed E Hassan. An empirical study of goto in c code from github repositories. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, pages 404–414. ACM, 2015.

48. Kapil Agrawal, Sadika Amreen, and Audris Mockus. Commit quality in five high performance computing projects. In *The 2015 International Workshop on Software Engineering for High Performance Computing in Science*, May 2015.

49. Qimu Zheng, Audris Mockus, and Minghui Zhou. A method to identify and correct problematic software activity data: Exploiting capacity constraints and data redundancies. In *ESEC/FSE'15*, Florence, Italy, 2015. ACM.

50. Minghui Zhou and Audris Mockus. Mining micro-practices from operational data. In *The 22nd ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, Hong Kong, November 16-21 2014.

51. Anh Nguyen Duc, Audris Mockus, Randy Hackbarth, and John Palframan. Forking and coordination in multi-platform development: a case study. In *ESEM*, Torino, Italy, September 2014.

52. Jiaxin Zhu, Minghui Zhou, and Audris Mockus. The relationship between folder use and the number of forks: A case study on github repositories. In *ESEM*, Torino, Italy, September 2014.

53. Feng Zhang, Audris Mockus, Iman Keivanloo, and Ying Zou. Towards building a universal defect prediction model. In *11th IEEE Working Conference on Mining Software Repositories*, May 30–31 2014.

54. Jialiang Xie, Qimu Zhengand, Minghui Zhou, and Audris Mockus. Product assignment recommender. In *ICSE'14 Demonstrations*, 2014.

55. Audris Mockus. Engineering big data solutions. In *ICSE'14 FOSE*, 2014.

56. Shane McIntosh, Martin Poehlmann, Elmar Juergens, Audris Mockus, Bram Adams, Ahmed E. Hassan, Brigitte Haupt, and Christian Wagner. Collecting and leveraging a benchmark of build system clones to aid in quality assessments. In *ICSE'14 Software Engineering In Practice*, 2014.

57. Jialiang Xie, Minghui Zhou, and Audris Mockus. Impact of triage: a study of Mozilla and Gnome. In *ESEM'13*, 2013.

58. Audris Mockus, Randy Hackbarth, and John Palframan. Risky files: An approach to focus quality improvement effort. In *9th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, 2013.

59. David Weiss and Audris Mockus. The chunking pattern. In *International Workshop on Data Analysis Patterns in Software Engineering*, 2013.

60. Feng Zhang, Audris Mockus, Ying Zou, Foutse Khomh, and Ahmed E. Hassan. How does context affect the distribution of software maintainability metrics? In *Proceedings of the 29th IEEE International Conference on Software Maintenance*, ICSM '13, 2013.

61. Dag I.K. Sjøberg, Bente Anda, and Audris Mockus. Questioning software maintenance metrics: a comparative case study. In *Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement*, ESEM'12, pages 107–110, New York, NY, USA, 2012. ACM.

62. Minghui Zhou and Audris Mockus. What make long term contributors: Willingness and opportunity in OSS community. In *ICSE 2012*, pages 518–528, Zürich, Switzerland, 2012.

63. Jialiang Xie, Audris Mockus, and Minghui Zhou. Visualizing evolution of software issue-tracking practices. In *ESEM'11*, Banff, Canada, Sep 22–23 2011.

64. Emad Shihab, Audris Mockus, Yasutaka Kamei, Bram Adams, and Ahmed E. Hassan. High-impact defects: a study of breakage and surprise defects. In *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*, ESEC/FSE '11, pages 300–310, New York, NY, USA, 2011. ACM.

65. Minghui Zhou and Audris Mockus. Does the initial environment impact the future of developers? In *ICSE 2011*, pages 271–280, Honolulu, Hawaii, May 21–28 2011.

66. Minghui Zhou and Audris Mockus. Growth of newcomer competence: Challenges of globalization. In *FSE/SDP Workshop on the Future of Software Engineering Research*, pages 442–447, Santa Fe, New Mexico, November 7–8 2010.

67. Audris Mockus. Organizational volatility and its effects on software defects. In *ACM SIGSOFT / FSE*, pages 117–126, Santa Fe, New Mexico, November 7–11 2010.

68. Minghui Zhou and Audris Mockus. Developer fluency: Achieving true mastery in software projects. In *ACM SIGSOFT / FSE*, pages 137–146, Santa Fe, New Mexico, November 7–11 2010.

69. Audris Mockus, Nachiappan Nagappan, and T Dinh-Trong, Trung. Test coverage and post-verification defects: A multiple case study. In *International Conference on Empirical Software Engineering and Measurement*, Lake Buena Vista, Florida USA, October 2009. ACM.

70. Audris Mockus. Amassing and indexing a large sample of version control systems: towards the census of public source code history. In *6th IEEE Working Conference on Mining Software Repositories*, May 16–17 2009.

71. Audris Mockus. Succession: Measuring transfer of code and developer productivity. In *2009 International Conference on Software Engineering*, Vancouver, CA, May 12–22 2009. ACM Press.

72. Minghui Zhou, Audris Mockus, and David Weiss. Learning in offshored and legacy software projects: How product structure shapes organization. In *ICSE Workshop on Socio-Technical Congruence*, Vancouver, Canada, May 19 2009.

73. Audris Mockus. Organizational volatility and developer productivity. In *ICSE Workshop on Socio-Technical Congruence*, Vancouver, Canada, May 19 2009.

74. Audris Mockus and David Weiss. Interval quality: Relating customer-perceived quality to process quality. In *2008 International Conference on Software Engineering*, pages 733–740, Leipzig, Germany, May 10–18 2008. ACM Press.

75. Hung-Fu Chang and Audris Mockus. Evaluation of source code copy detection methods on FreeBSD. In *5th Working Conference on Mining Software Repositories*. ACM Press, May 10–11 2008.

76. Audris Mockus. Large-scale code reuse in open source software. In *ICSE'07 Intl. Workshop on Emerging Trends in FLOSS Research and Development*, Minneapolis, Minnesota, May 21 2007.

77. Sunghun Kim, Thomas Zimmermann, Miryung Kim, Ahmed E. Hassan, Audris Mockus, Tudor Gîrba, Martin Pinzger, E. James Whitehead Jr., and Andreas Zeller. Ta-re: an exchange language for mining software repositories. In *ICSE'06 Workshop on Mining Software Repositories*, pages 22–25, Shanghai, China, May 22-23 2006.

78. Hung-Fu Chang and Audris Mockus. Constructing universal version history. In *ICSE'06 Workshop on Mining Software Repositories*, pages 76–79, Shanghai, China, May 22-23 2006.

79. Audris Mockus. Empirical estimates of software availability of deployed systems. In *2006 International Symposium on Empirical Software Engineering*, pages 222–231, Rio de Janeiro, Brazil, September 21-22 2006. ACM Press.

80. Birgit Geppert, Audris Mockus, and Frank Rößler. Refactoring for changeability: A way to go? In *Metrics 2005: 11th International Symposium on Software Metrics*, Como, September 2005. IEEE CS Press.

81. Audris Mockus, Ping Zhang, and Paul Li. Drivers for customer perceived software quality. In *ICSE 2005*, pages 225–233, St Louis, Missouri, May 2005. ACM Press.

82. Audris Mockus. Analogy based prediction of work item flow in software projects: a case study. In *2003 International Symposium on Empirical Software Engineering*, pages 110–119, Rome, Italy, October 2003. ACM Press.

83. James Herbsleb and Audris Mockus. Formulation and preliminary test of an empirical theory of coordination in software engineering. In *2003 International Conference on Foundations of Software Engineering*, Helsinki, Finland, October 2003. ACM Press.

84. Audris Mockus, David M. Weiss, and Ping Zhang. Understanding and predicting effort in software projects. In *2003 International Conference on Software Engineering*, pages 274–284, Portland, Oregon, May 3-10 2003. ACM Press.

85. Audris Mockus and James Herbsleb. Expertise browser: A quantitative approach to identifying expertise. In *2002 International Conference on Software Engineering*, pages 503–512, Orlando, Florida, May 19-25 2002. ACM Press.

86. Stacie Hibino and Audris Mockus. handiMessenger: Awareness-enhanced universal communication for mobile users. In Fabio Paternò, editor, *Mobile Human-Computer Interaction, 4th International Symposium, Mobile HCI 2002, Pisa, Italy, September 18-20, 2002, Proceedings*, volume 2411 of *Lecture Notes in Computer Science*, pages 170–183, Pisa, Italy, September, 18-20 2002. Springer.

87. J. A. Espinosa, R.E. Kraut, S. A. Slaughter, J. F. Lerch, J. D. Herbsleb, and A. Mockus. Shared mental models, familiarity, and coordination: A multi-method study of distributed software teams. In *Proc. International Conference in Information Systems*, pages 513–518, Barcelona, Spain, December 15–18 2001.

88. James D. Herbsleb, Audris Mockus, Thomas A. Finholt, and Rebecca E. Grinter. An empirical study of global software development: Distance and speed. In *23nd International Conference on Software Engineering*, pages 81–90, Toronto, Canada, May 12-19 2001.

89. A Mockus and JD Herbsleb. Challenges of global software development. In *Metrics 2001: Seventh International Symposium on Software Metrics*, pages 182–184, London, England, April 4-6 2001.

90. HP Siy, A Mockus, JD Herbsleb, M Krishnan, and GT Tucker. Making the software factory work: Lessons from a decade of experience. In *Metrics 2001: Seventh International Symposium on Software Metrics*, pages 317–327, London, England, April 4-6 2001.

91. James D. Herbsleb, Audris Mockus, Thomas A. Finholt, and Rebecca E. Grinter. Distance, dependencies, and delay in a global collaboration. In *ACM 2000 Conference on Computer Supported Cooperative Work*, pages 319–328, Philadelphia, Pennsylvania, USA, December 2–6 2000. ACM, ACM Press.

92. Audris Mockus and Lawrence G. Votta. Identifying reasons for software change using historic databases. In *International Conference on Software Maintenance*, pages 120–130, San Jose, California, October 11-14 2000.

93. A. Mockus, R. F. Fielding, and J. Herbsleb. A case study of open source development: The Apache server. In *22nd International Conference on Software Engineering*, pages 263–272, Limerick, Ireland, June 4-11 2000.

94. S.L. Hibino, T. Graves, and A. Mockus. A web based approach to interactive visualization in context. In *Advanced Visual Interfaces*, pages 181–188, Palermo, Italy, May 23-26 2000.

95. Harvey Siy and Audris Mockus. Measuring domain engineering effects on software coding cost. In *Metrics 99: Sixth International Symposium on Software Metrics*, pages 304–311, Boca Raton, Florida, November 4–6 1999.

96. Audris Mockus. Analysis of software evolution data. In *Joint Statistical Meetings, 1999. Abstracts*, page 121, Baltimore, Maryland, August 8-14 1999.

97. D. Atkins, T. Ball, T. Graves, and A. Mockus. Using version control data to evaluate the effectiveness of software tools. In *1999 International Conference on Software Engineering*, pages 324–333. ACM Press, May 16–22 1999.

98. K.C. Cox, S.L. Hibino, L. Hong, A. Mockus, and G.J. Wills. Infostill: A task-oriented framework for analyzing data through information visualization. In *Proc. IEEE Information Visualization Late Breaking Hot Topics*, pages 19–22, San Francisco, CA, October 1999.

99. Nancy Staudenmayer, Todd Graves, J. Steve Marron, Audris Mockus, Dewayne Perry, Harvey Siy, and Lawrence Votta. Adapting to a new environment: How a legacy software organization copes with volatility and change. In *5th International Product Development Management Conference*, Como, Italy, May 1998.

100. Audris Mockus. Navigating aggregation spaces. In *IEEE Information Visualization Symposium 1998 Late Breaking Hot Topics Proceedings*, pages 29–32, Research Triangle Park, NC, 1998. IEEE Computer Society Press.

101. Todd L. Graves and Audris Mockus. Inferring programmer effort from software databases. In *22nd European Meeting of Statisticians and 7th Vilnius Conference on Probability Theory and Mathematical Statistics*, page 334, Vilnius, Lithuania, August 1998.

102. Todd L. Graves and Audris Mockus. Inferring change effort from configuration management data. In *Metrics 98: Fifth International Symposium on Software Metrics*, pages 267–273, Bethesda, Maryland, November 1998.

103. Stephen G. Eick, Audris Mockus, Todd L. Graves, and Alan F. Karr. Web-based text visualization. In Wolfgang Bandilla and Frank Faulbaum, editors, *SoftStat '97 Advances in Statistical Software 6*, pages 3–10. Lucius & Lucius, March 1997.

104. W.F. Eddy, M. Fitzgerald, C.R. Genovese, A. Mockus, and D.C. Noll. Functional imaging analysis software - computational olio. In A. Prat, editor, *Proceedings in Computational Statistics*, pages 39–49. Physica-Verlag, Heidelberg, 1996.

105. A. Mockus, W.F. Eddy, S.Y. Chang, and K.R. Thulborn. Software for the visualization of fMRI data. In *Proceedings of the International Society for Magnetic Resonance in Medicine Fourth Scientific Meeting and Exhibitionn*, page 1774, 1996.

106. W.F Eddy, M. Fitzgerald, C. Genovese, and A. Mockus. The challenge of functional magnetic resonance imaging. In *Massive Data Sets*, pages 39–45, Washington, D.C., 1996. National Academy Press.

107. W. F. Eddy and A. Mockus. Discovering, describing, and understanding spatial-temporal patterns of disease using dynamic graphics. In *Proceedings of the 25th Public Health Conference on Records and Statistics, U.S. Department of Health and Human Services*, pages 14–17, 1995.

108. A. Mockus, J. Mockus, and L. Mockus. Discrete optimization, information based complexity, and Bayesian heuristics approach. In *Proceedings of International Symposium on Operations Research with Applications in Engineering Technology and Management (ISORA'95), August 19-22*, Beijing, China, 1995.

109. A. Mockus. Estimating dependencies from spatial averages. In *American Statistical Association: 1995 Proceedings of the Section on Statistical Graphics*, Orlando, Florida, August 1995.

110. W.F. Eddy and A. Mockus. Visualizing spatial and temporal trends for the incidence of six diseases. In *Symposium on Statistical Methods 1995. Small Area Statistics in Public Health*, Atlanta, Georgia, 1995.

111. M. Behrmann, P.A. Carpenter, S.Y. Chang, W.F. Eddy, J.S. Gillen, M.A. Just, T.A. Keller, A. Mockus, T.A. Tasciyan, and K.R. Thulborn. Test-retest reproducibility during fMRI studies: Primary visual and cognitive paradigms. In *Proceedings of the Society of Magnetic Resonance, Third Scientific Meeting*, volume 2, page 843, 1995.

112. W.F. Eddy and A. Mockus. An interactive image index. In *American Statistical Association: 1994 Proceedings of the Section on Statistical Graphics*, Toronto, Canada, 1994.

113. A. Mockus, J. Mockus, and L. Mockus. Bayesian approach in stochastic and heuristic methods of global and discrete optimization. In *Abstracts, Second World Meeting, International Society for Bayesian Analysis*, pages 10–11, Alicante, Spain, 1994.

114. W.F. Eddy and A. Mockus. Noninteractive display of functions with temporal and spatial variation (or surreal-time graphics). In *New Directions in Statistical Data Analysis and Robustness*, pages 47–60. Birkhäuser Verlag, Basel, 1993.

115. W.F. Eddy, M.M. Meyer, A. Mockus, M.J. Schervish, K. Tan, and K. Viele. Smoothing census adjustment factors - an application of high performance computing. In *Computing Science and Statistics: Proceedings of the 24th Symposium on the Interface*, pages 503–510, 1992.

Articles in Refereed Journals or Book Chapters

116. Anna Maria Nau, Philipp Ditto, Dawnie Steadman, and Audris Mockus. Identifying factors that help improve existing decomposition-based pmi estimation methods. *Journal of Forensic Anthropology*, 2025.

117. Anna-Maria Nau, Phillip Ditto, Dawnie Wolfe Steadman, and Audris Mockus. Towards automation of human stage of decay identification: An artificial intelligence approach. *Journal of Forensic and Legal Medicine*, 116:102987, 2025.

118. Mahmoud Jahanshahi, David Reid, and Audris Mockus. Beyond dependencies: The role of copy-based reuse in open source software development. *ACM Transactions on Software Engineering and Methodology*, 2025.

119. P.C. Rigby, S. Rogers, S. Saleem, P. Suresh, D. Suskin, P. Riggs, C. Maddila, N. Nagappan, and A. Mockus. Improving code reviewer recommendation: Accuracy, latency, workload, and bystanders. *ACM Transactions on Software Engineering and Methodology*, 2025.

120. Anna-Maria Nau, Sara Mousavi, Dylan Lee, Rayhan Hossain, Tatianna Griffin, Dawnie Wolfe Steadman, and Audris Mockus. Icputrd: Image cloud platform for use in tagging and research on decomposition. *Journal of Forensic Sciences*, 2023.

121. Audris Mockus, Rui Abreu, Peter Rigby, David Amsallem, Parveen Bansal, Kaavya Chinniah, Brian Ellis, Peng Fan, Jun Ge, Bingjie He, Kelly Hirano, Sahil Kumar, Megh Mehta, Ajay Lingapuram, Andrew Loe, Venus Montes, Maher Saba, Gursharan Singh, Matt Steiner, Weiyan Sun, Siri Uppalapati, and Nachiappan Nagappan. Leveraging risk models to improve productivity for effective code un-freeze at scale. *ACM Transactions on Software Engineering and Methodology*, 2025.

122. Kai Gao, Zhixing Wang, Audris Mockus, and Minghui Zhou. On the variability of software engineering needs for deep learning: Stages, trends, and application types. *IEEE Transactions on Software Engineering*, 2022. ISSN: 0098-5589 DOI:10.1109/TSE.2022.3163576.

123. Yuxing Ma, Audris Mockus, Russell Zaretzki, Bogdan Bichescu, and Randy Bradley. A methodology for analyzing uptake of softwaretechnologies among developers. *IEEE Transactions on Software Engineering*, 2020.

124. Peter C. Rigby Andrey Krutauz, Tapajit Dey and Audris Mockus. Do code review measures explain the incidence of post-release defects? *International Journal of Empirical Software Engineering*, 25:3323–3356, 2020.

125. Yuxing Ma, Tapajit Dey, Chris Bogart, Sadika Amreen, Marat Valiev, Adam Tutko, David Kennard, Russell Zaretzki, and Audris Mockus. World of code: Enabling a research workflow for mining and analyzing the universe of open source vcs data. *International Journal of Empirical Software Engineering*, 2020.

126. Tapajit Dey and Audris Mockus. Deriving a usage-independent software quality metric. *International Journal of Empirical Software Engineering*, 25:1596–1641, 2020.

127. Randy V. Bradley, Audris Mockus, Yuxing Ma, and Russell Zaretzkiand Bogdan C. Bichescu. Coordinating interdependencies in an open source software project: A replication of lindberg, et al. *AIS Transactions on Replication Research*, 2020. DOI: 10.17705/1atrr.00057.

128. Sadika Amreen, Yuxia Zang, Chris Bogart, Russell Zaretzki, and Audris Mockus. Alfaa: Active learning fingerprint based anti-aliasing for correcting developer identity errors in version control systems. *International Journal of Empirical Software Engineering*, 2019.

129. Yuxia Zhang, Minghui Zhou, Audris Mockus, and Zhi Jin. Companieśparticipation in oss development - an empirical study of openstack. *IEEE Transactions on Software Engineering*, September 2019.

130. Sadika Amreen, Bogdan Bichescu, Randy Bradley, Tapajit Dey, Yuxing Ma, Audris Mockus, Sara Mousavi, and Russell Zaretzki. A methodology for measuring floss ecosystems. In *Towards Engineering Free/Libre Open Source Software (FLOSS) Ecosystems for Impact and Sustainability: Communications of NII Shonan Meetings*. Springer, 2019.

131. Introduction. In *Towards Engineering Free/Libre Open Source Software (FLOSS) Ecosystems for Impact and Sustainability: Communications of NII Shonan Meetings*. 2019.

132. Feng Zhang, Audris Mockus, Iman Keivanloo, and Ying Zou. Towards building a universal defect prediction model with rank transformed predictors. *International Journal of Empirical Software Engineering*, 21(5):2107–2145, 2016.

133. Minghui Zhou, Audris Mockus, Xiujuan Ma, Lu Zhang, and Hong Mei. Inflow and retention in oss communities with commercial involvement: A case study of three hybrid projects. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 25(2):13, 2016.

134. A Mockus. Operational data are missing, incorrect, and decontextualized. In *Perspectives on Data Science for Software Engineering*, pages 317–322. Elsevier, 2016.

135. Randy Hackbarth, Audris Mockus, John Palframan, and David Weiss. Assessing the state of software in a large enterprise: A 12-year retrospective. In *The Art and Science of Analyzing Software Data*, pages 411–451. Elsevier, 2016.

136. Ian Gorton, Ayse Basar Bener, and Audris Mockus. Software engineering for big data systems. *IEEE Software*, 33(2):32–35, 2016.

137. R. Hackbarth, A. Mockus, J. Palframan, and R. Sethi. Customer quality improvement of software systems. *Software, IEEE*, 33(4):40–45, 2016.

138. Minghui Zhou and Audris Mockus. {Who Will Stay in the FLOSS Community? Modelling Participant's Initial Behaviour}. *IEEE Transactions on Software Engineering*, 41(1):82–99, Jan 2015.

139. Shane McIntosh, Meiyappan Nagappan, Bram Adams, Audris Mockus, and Ahmed E. Hassan. A large-scale empirical study of the relationship between build technology and build maintenance. *Journal of Empirical Software Engineering*, 2014.

140. Yasutaka Kamei, Emad Shihab, Bram Adams, Ahmed E. Hassan, Audris Mockus, Anand Sinha, and Naoyasu Ubayashi. A large-scale empirical study of just-in-time quality assurance. *IEEE Transactions on Software Engineering*, 2013.

141. Dag I.K. Sjoberg, Aiko Yamashita, Bente Anda, Audris Mockus, and Tore Dyba. Quantifying the effect of code smells on maintenance effort. *IEEE Transactions on Software Engineering*, 2013.

142. Audris Mockus and David Weiss. Software chunks and distributed development. In Christof Ebert, editor, *Global Software and IT: a Guide to Distributed Development, Projects, and Outsourcing*, pages 69–82. Willey, 2011.

143. Randy Hackbarth, Audris Mockus, John Palframan, and David Weiss. Assessing the state of software in a large enterprise. *Journal of Empirical Software Engineering*, 10(3):219–249, 2010.

144. Marcelo Cataldo, Audris Mockus, Jeffrey A. Roberts, and James D. Herbsleb. Software dependencies, the structure of work dependencies and their impact on failures. *IEEE Transactions on Software Engineering*, 2009.

145. Bente C.D. Anda, Dag I.K. Sjøberg, and Audris Mockus. Variability and reproducibility in software engineering: A study of four companies that developed the same system. *IEEE Transactions on Software Engineering*, 35(3), May/June 2009.

146. Audris Mockus. Missing data in software engineering. In J. Singer et al., editor, *Guide to Advanced Empirical Software Engineering*, pages 185–200. Springer-Verlag, 2008.

147. Audris Mockus. Software support tools and experimental work. In V Basili and et al, editors, *Empirical Software Engineering Issues: Critical Assessments and Future Directions*, volume LNCS 4336, pages 91–99. Springer, 2007.

148. Audris Mockus, Roy T. Fielding, and James Herbsleb. Two case studies of open source software development: Apache and Mozilla. In J Feller and et al, editors, *Perspectives on Free and Open Source Software*, pages 163—210. MIT Press, 2005.

149. D. L. Atkins, A. Mockus, and H. P. Siy. *Value Based Software Engineering*, chapter Quantifying the Value of New Technologies for Software Development, pages 327–344. Springer Verlag Berlin Heidelberg, 2006.

150. J. D. Herbsleb and A. Mockus. An empirical study of speed and communication in globally-distributed software development. *IEEE Transactions on Software Engineering*, 29(6):481–494, June 2003.

151. D. Atkins, T. Ball, T. Graves, and A. Mockus. Using version control data to evaluate the impact of software tools: A case study of the version editor. *IEEE Transactions on Software Engineering*, 28(7):625–637, July 2002.

152. Audris Mockus, Roy T. Fielding, and James Herbsleb. Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology*, 11(3):1–38, July 2002.

153. S.G. Eick, T.L. Graves, A.F. Karr, A. Mockus, and P. Shuster. Visualizing software changes. *IEEE Transactions on Software Engineering*, 28(4):396–412, April 2002.

154. Audris Mockus and David M. Weiss. Globalization by chunking: a quantitative approach. *IEEE Software*, 18(2):30–37, March 2001.

155. T. Graves and A. Mockus. Identifying productivity drivers by modeling work units using partial data. *Technometrics*, 43(2):168–179, May 2001.

156. Stephen G. Eick, Todd L. Graves, Alan F. Karr, J. S. Marron, and Audris Mockus. Does code decay? assessing the evidence from change management data. *IEEE Transactions on Software Engineering*, 27(1):1–12, 2001.

157. Audris Mockus and David M. Weiss. Predicting risk of software changes. *Bell Labs Technical Journal*, 5(2):169–180, April–June 2000.

158. D. Atkins, A. Mockus, and H. Siy. Measuring technology effects on software change cost. *Bell Labs Technical Journal*, 5(2):7–18, April–June 2000.

159. W.F Eddy, M. Fitzgerald, C. Genovese, N. Lazar, A. Mockus, and Welling J. The challenge of functional magnetic resonance imaging. *Journal of Computational and Graphical Statistics*, 8(3):545–558, September 1999.

160. J. Mockus, A. Mockus, and L. Mockus. Bayesian approach for randomization of heuristic algorithms in discrete programming. In Panos Pardalos, Sanguthevar Rajasekaran, and Jose Rolim, editors, *DIMACS series in discrete mathematics and theoretical computer science*, volume 43, pages 161–177. American Mathematical Society, 1999.

161. Audris Mockus. Estimating dependencies from spatial averages. *Journal of Computational and Graphical Statistics*, 7(4):501–513, 12 1998.

162. Stephen G. Eick, Audris Mockus, Todd L. Graves, and Alan F. Karr. A web laboratory for software data analysis. *World Wide Web*, 1(2):55–60, 1998.

163. Audris Mockus, Adam Porter, Harvey Siy, and Lawrence G. Votta. Understanding the sources of variation in software inspections. *ACM Transactions on Software Engineering and Methodology*, 7(1), January 1998.

164. A. Mockus, J. Mockus, and L.Mockus. Bayesian heuristic approach (BHA) and applications to discrete optimization. *Fields Institute Communications*, 18:153–165, 1998.

165. W.F. Eddy and A. Mockus. An interactive icon index: Images of the outer planets. *Journal of Computational and Graphical Statistics*, 5(1):100–111, 1996.

166. J. Mockus, W.F. Eddy, A. Mockus, L. Mockus, and G. Reklaitis. *Bayesian Heuristic Approach to Discrete and Global Optimization*. Kluwer Academic Publishers, Dordrecht, 1997.

167. W. F. Eddy and A. Mockus. Dynamic visualization in modelling and optimization of ill defined problems. In C.A. Floudas and P.M. Pardalos, editors, *State of the Art in Global Optimization: computational methods and applications*, pages 499–520. Kluwer Academic Publishers, Dordrecht, 1996.

168. W.F. Eddy, A. Mockus, and S. Oue. Approximate single linkage cluster analysis of large data sets in spaces of high dimension. *Computational Statistics and Data Analysis*, 23:29–43, 1996.

169. M. Lavine and A. Mockus. A nonparametric Bayes method for isotonic regression. *Journal of Statistical Planning and Inference*, 46:235–248, 1995.

170. W.F. Eddy and A. Mockus. An example of the estimation and display of a smoothly varying function of time and space - the incidence of mumps disease. *Journal of the American Society for Information Science*, 45(9):686–693, 1994.

171. A. Mockus, J. Mockus, and L. Mockus. Adapting stochastic and heuristic methods for discrete optimization problems. *Informatica*, 5(1):123–166, 1994.

172. W.F. Eddy and A. Mockus. An example of noninteractive dynamic graphics for the visualization of manufacturing process data. *International Statistical Review*, 61(1):81–95, 1993.

173. A. Mockus and L. Mockus. Designing software for global optimization. *Informatica*, 1(1):71–88, 1990.

Tutorials and Talks

174. Audris Mockus. Securing large language model software supply chains, September 13 2023. ASE'23 LLMs in Software Engineering.

175. Audris Mockus. World of code: Enabling a research workflow for mining and analyzing the universe of open source vcs data, December 3 2021. DAMSS'21 Keynote.

176. Audris Mockus. Insights from open source supply chains, August 30 2019. FSE'19 Industry Keynote.

177. Audris Mockus. Knowledge flows in open source software supply chains, Nov 4 2017. Keynote at the 16th National Software Applicatio Conference.

178. Bian Fitzgerald, Audris Mockus, and Minghui Zhou. Towards engineering free/libre open source software (floss) ecosystems for impact and sustainability, June 2017.

179. Audris Mockus. Comparative linguistics meta-meta-science and knowledge flows, May 2017. Panel at International Conference on Global Software Engineering.

180. Audris Mockus. Evidence engineering. Invited talk at Business Analytics Forum.

181. Audris Mockus. Towards evidence engineering. Invited talk at Concordia University.

182. Audris Mockus. Evidence engineering, February 2015. Keynote at 8th India Software Engineering Conference.

183. Audris Mockus. Defect prediction and software risk, June 30 2014. Promise'14 Keynote/Tutorial.

184. Audris Mockus. Is mining software repositories data science?, June 30 2014. Talk at Vilnius University.

185. Audris Mockus. Operational data are not experimental data, June 22-27 2014. Dagstuhl Seminar 14261: Software Development Analytics.

186. Audris Mockus. Is mining software repositories data science?, May 2014. Keynote at 11th IEEE Working Conference on Mining Software Repositories 2014.

187. A. Mockus, M. Nagappan, and A. Hassan. Tutorial on statistics in software engineering: Pitfalls and good practices, August 2013. 25-30 attendees.

188. Audris Mockus. The paradox of software quality, November 2013. Distinguished Seminar, Queen's University, Kingston, Canada.

189. Audris Mockus. Law of minor release: More bugs $\Longrightarrow$ better software quality. In *International Workshop on Principles of Software Evolution*, St Petersburg, Russia, Aug 18-19 2013. Keynote.

190. Audris Mockus. Quantifying and conveying risk in software development. In *1st Symposium on Mining Software Archives*, Monte Verita, Switzerland, March 10-15 2013. Keynote.

191. James Herbsleb, Marcelo Cataldo, Daniela Damian, Premkumar Devenbu, Steve Easterbrook, and Audris Mockus. Socio-technical congruence (stc 2008). In *Companion of the 30th international conference on Software engineering*, 2008.

192. Ahmed E. Hassan, Richard C. Holt, and Audris Mockus. Report on MSR 2004: International workshop on mining software repositories. In *ACM SIGSOFT Software Engineering Notes*, 2005.

193. Audris Mockus. Measuring distributed software engineering. In *International Conference on Global Software Engineering*, Princeton, New Jersey, August 25, 2010 2010. Keynote.

194. Audris Mockus. Globalization and transfer of ownership. In *International Conference on Global Software Engineering*, Princeton, New Jersey, August 25, 2010 2010. Panel: Impact of Future Communication Technologies on GSD.

195. Audris Mockus. Measurement in science and software engineering. In *1st Symposium on Mining Software Archives*, Monte Verita, Switzerland, March 17 2010. Keynote.

196. Audris Mockus. Globalization and the future developer. In *New Frontiers for Empirical Software Engineering*, Dagstuhl, Germany, March, 2010 2010. Panel: Future of Globally Distributed Software Development.

197. Audris Mockus. Using software changes to understand software projects, February 2010. Distinguished Seminar, Queen's University, Kingston, Canada.

198. Audris Mockus. Towards understanding of software changes, 2009. Peking University, Beijing.

199. Audris Mockus. Using software changes to understand software projects, 2009. Institute of Software, Chinese Academy of Sciences, Beijing.

200. Audris Mockus. Domain-specific defect models. In *International Workshop on Defects in Large Software System*, Seattle, WA, July 20 2008. Keynote.

201. Audris Mockus. Large scale reuse in open source software, 2008. Institute of Mathematics and Informatics, Vilnius.

202. Audris Mockus. Large scale reuse in open source software, 2008. Kaunas University of Technology, Kaunas.

203. Audris Mockus. Faces of software quality, 2008. Microsoft Research, Redmond.

204. Audris Mockus and David Weiss. Industrial strength software measurement, 2006. Peking University, Beijing.

205. Audris Mockus. Software changes and software engineering. In *Multi-Version Program Analysis*. Schloss Dagstuhl, Germany, 26 June 2005. Keynote.

206. Audris Mockus and David Weiss. Software changes and software engineering, 2005. Brooklyn Poly, Brooklyn, NY.

207. Audris Mockus. Effects of distributed software development and virtual teams, 2004. University of Victoria, Victoria.

208. Audris Mockus. Effects of distributed software development and virtual teams, 2004. Simula Labs, Oslo.

209. Audris Mockus. Software changes: from insights to solutions, 2002. Carnegie Mellon University, Pittsburgh.

210. Audris Mockus. Analysis of software changes, 2002. University of Maryland, College Park.

211. Audris Mockus. Expertise browser. In *Joint Statistical Meetings, 2001*, Atlanta, Georgia, August 5-9 2001.

212. Audris Mockus. Role of change history in empirical studies of software, 1998. University of Karlsruhe.

213. Audris Mockus. Modeling software evolution, 1998. Texas A&M University, College Station, TX.

214. Audris Mockus. If software could talk: why changes are made, 1998. Fraunhoffer IESE, Kaiserslautern.

215. Audris Mockus. Inferring programmer effort from software change data, 1998. Institute of Mathematics and Informatics, Vilnius.

Workshops

216. Adam Tutko, Austin Z Henley, and Audris Mockus. How are software repositories mined? a systematic literature review of workflows, methodologies, reproducibility, and tools. *arXiv preprint arXiv:2204.08108*, 2022.

217. Audris Mockus. Tutorial: Open source software supply chains. In *India Software Engineering Conference*, 2022.

218. Audris Mockus. Woc tutorial, 2019.

219. Audris Mockus. Oscar: Open source supply chains: Asesment and mitigation of risks, 2017.

220. Tapajit Dey and Audris Mockus. A Matching Based Theoretical Framework for Estimating Probability of Causation. *arXiv e-prints*, page arXiv:1808.04139, Aug 2018.

221. James Herbsleb, Marcelo Cataldo, Daniela Damian, Premkumar Devanbu, Steve Easterbrook, and Audris Mockus. Socio-technical congruence (STC 2008). In *ICSE Companion '08: Companion of the 30th international conference on Software engineering*, pages 1027–1028, New York, NY, USA, 2008. ACM.

222. Audris Mockus. How to run empirical studies using project repositories. 4th International Advanced School of Empirical Software Engineering, September 20, 2006, Rio de Janeiro, Brazil 2006. Tutorial.

223. Daniel German and Audris Mockus. Automating the measurement of open source projects. In *ICSE '03 Workshop on Open Source Software Engineering*, page Automating the Measurement of Open Source Projects, Portland, Oregon, May 3-10 2003.

224. Audris Mockus. Measurement in software projects: taking advantage of version control repositories. In *International Software Engineering Network, 2002*, Nara, Japan, October 2002.

225. Audris Mockus and James Herbsleb. Why not improve coordination in distributed software development by stealing good ideas from open source. In *ICSE '02 Workshop on Open Source Software Engineering*, pages 35–37, Orlando, FL, May 2002.

226. Audris Mockus T. Ball, Stephen G. Eick. Web-based analysis of large-scale software systems. In *ICSE '97 Workshop on Software Engineering (on) the World Wide Web*, Boston, MA, May 1997.

227. Nancy Staudenmayer, Todd Graves, and Audris Mockus. Adapting to a new environment: How a legacy software organization copes with volatility and change. In *Academy of Management Chicago 1999 Conference*, Chicago, Illinois, August 1999.

228. Audris Mockus, Todd L. Graves, and Alan F. Karr. Modelling software changes. In C.E. Minder and H. Friedl, editors, *Good Statistical Practice*, pages 175–179. Austrian Statistical Society, Wien, Austria, July 1997. Proceedings of the 12th International Workshop on Statistical Modeling, Biel/Bienne.

229. William F. Eddy and Audris Mockus Incidence of Disease Mumps ASA Video Library (1994) 15:00 minutes http://www.bell-labs.com/topic/societies/asagraphics/library/

230. William F. Eddy and Audris Mockus Manufacturing Process Data ASA Video Library (1993) 9:00 minutes http://www.bell-labs.com/topic/societies/asagraphics/library/

Technical Reports

231. Audris Mockus. Succession: Measuring transfer of code and developer productivity. Technical report, Avaya Labs research, 2008.

232. Audris Mockus. Transfer of code ownership, implicit teams, and organizational tomography. Technical Report ALR-2008-010, CID 135147, Avaya Labs research, 2008.

233. M Ardis, KC Cox, SL Hibino, L Hong, A Mockus, and GJ Wills. Building information visualizations: A commonality analysis. Technical report, Lucent Technologies, 2000.

234. A Mockus and DM Weiss. Globalization by chunking: A quantitative approach. Technical Report 10009677-000727-03TM, Lucent Technologies, 2000.

235. JD Herbsleb, M Krishnan, A Mockus, HP Siy, and GT Tucker. Making the software factory work: Lessons from a decade of experience. Technical Report 10009677-000404-02TM, Lucent Technologies, 2000.

236. A Mockus, H Siy, T Sundresh, J Chen, and TL Graves. Role of change size, complexity, and developer expertise in predicting the quality of a software update. Technical Report 10009677-000324-01TM, Lucent Technologies, 2000.

237. KC Cox, SL Hibino, L Hong, A Mockus, and GJ Wills. A method for graphically displaying an overview of multiple database tables. Technical Report 10009677-991208-07TM, Lucent Technologies, 2000.

238. S Eick, P Schuster, A Mockus, T Graves, and A Karr. Visualizing software changes. Technical Report BL0113590-990827-13TM, Lucent Technologies, 1999.

239. A Mockus and T Graves. Identifying productivity drivers by modeling work units using partial data. Technical Report BL0113590-990513-09TM, Lucent Technologies, 1999.

240. A. Mockus, S. G. Eick, T. L. Graves, and A. F. Karr. On measurement and analysis of software changes. Technical Report BL0113590-990401-06TM, Bell Laboratories, Lucent Technologies, 1999.

241. T Graves, J Marron, A Mockus, D Perry, H Siy, N Staudenmayer, and L Votta. Adapting to a new environment: How a legacy software organization copes with volatility and change. Technical Report BL0113590-980630-07TM, Lucent Technologies, 1998.

242. A Mockus and LG Votta. Identifying reasons for software changes using historic databases. Technical Report BL0113590-980410-04, Lucent Technologies, 1998.

243. SG Eick, TL Graves, AF Karr, and A Mockus. Web-based text visualization. Technical Report BL0112590-961025-22TM, Lucent Technologies, 1997.

244. Audris Mockus, Adam A. Porter, Harvey P. Siy, and Lawrence G. Votta. Understanding the sources of variation in software inspections. Technical Report BL0112590-960416-12TM, Bell Laboratories, Lucent Technologies, Naperville, IL, April 1996.

245. A. Mockus. *Predicting a Space-Time Process from Aggregate Data Exemplified by the Animation of Mumps Disease*. PhD thesis, Carnegie Mellon University, 1994.

246. Sadika Amreen. *Methods of Disambiguating andDe-anonymizing Authorship in LargeScale Operational Data*. PhD thesis, 2019.

247. Yuxing Ma. *Software Supply Chain (SSC) Development and Application*. PhD thesis, 2020.

248. Tapajit Dey. *Modeling User-affected Software Properties for Open Source Software Supply Chains*. PhD thesis, 2020.

249. Sara Mousavi. *Auto-curation of Large Evolving Image Datasets*. PhD thesis, 2021.

250. David Reid. *Applying the Universal Version History Concept to Help De-Risk Copy-Based Code Reuse*. PhD thesis, 2023.

251. Anna Maria Nau. *Applications of Machine Learning in Forensic Investigations of Human Decomposition*. PhD thesis, 2024.

252. Mahmoud Jahnshahi. *Copy-Based Reuse and its Implications in Open Source Software Supply Chains*. PhD thesis, 2025.

Professional

University, college, department service

- Chair, graduate admissions comittee 2024 - present
- Chair, promotion and tenure comittee 2024 - present
- Faculty Senate 2016-present
- Bredesen Center Faculty, 2018 - present
- HITES11 Faculty Participant, Summer 2015
- EECS Graduate Committee Member, 2017 - present
- EECS Strategy Committee Member, 2014-2015

Community service

Committee Member, panel on Transparency and Reproducibility of Federal Statistics for the National Center for Science and Engineering Statistics, 2019-2020

Program Co-Chair, ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Oulu, Finland, 2018.

Co-Organizer, Shonan meeting: Towards engineering free/libre open source software (floss) ecosystems for impact and sustainability, Shaonan Village, Japan, June 26-29, 2017.

Program Co-Chair, ESEC-FSE'15 Industry, Bergamo, Italy 2015.

General Chair, MSR'10, Cape Town, South Africa 2010.

Program Co-Chair, Second International Conference on Global Software Engineering, Munich, 2007.

Program Co-Chair, Tenth International Symposium on Software Metrics, Chicago, 2004.

Co-Organizer, ICSE Workshop on Socio-Technical Congruence, 2008,2009

Co-Organizer, 1st International Workshop on Mining Software Repositories, Edinburgh, Scotland, 2004.

Editorial board, IEEE Transactions on Software Engineering 2005-2009 and 2013-2017

Editorial board, International Journal of Empirical Software Engineering 2002-.

Co-Editor, IEEE Software special issue on Software Engineering for Big Data Systems, 2016.

Co-Editor, IEEE Transactions on Software Engineering special issue on Socio Technical Congruence, 2009-2010.

Co-Editor, IEEE Transactions on Software Engineering special issue on Mining Software Repositories, 2004-2005.

Advisory committee, 1st International Workshop on Replication in Empirical Software Engineering Research, Cape Town, South Africa.

Member, Program Committee, Foundations of Software Engineering, Seattle, 2026

Member, Program Committee, International Conference on Software Engineering, 2023

Member, Program Committee, International Conference on Software Engineering, 2020

Member, Program Committee, Foundations of Software Engineering, Seattle, 2019, Tartu, Estonia

Member, Program Committee, International Conference on Software Engineering, Software Engineering in Practice, Montréal, Canada, 2019

Member, Program Committee, International Conference on Software Engineering, Software Engineering in Society, Montréal, Canada, 2019

Member, Program Committee, ACM / IEEE International Symposium on Empirical Software Engineer- ing and Measurement, 2019.

Member, Program Committee, 15th IEEE Working Conference on Mining Software Repositories, Montréal, Canada, 2019.

Member, Program Committee, 33rd IEEE/ACM International Conference on Automated Software Engineering, Montpellier, France, 2018

Member, Program Committee, ACM / IEEE International Symposium on Empirical Software Engineer- ing and Measurement, Toronto, 2017.

Member, Program Committee, 14th IEEE Working Conference on Mining Software Repositories, Buenos Aires, Argentina, 2017.

Member, Program Committee, Foundations of Software Engineering - Industry Track, Paderborn, Germany, 2017

Member, Program Committee, 12th International Conference on Global Software Engineering, Buenos Aires, Argentina, 2017.

Member, Program Committee, Foundations of Software Engineering, Seattle, 2016

Member, Program Committee, Foundations of Software Engineering - Industry Track, Seattle, 2016

Member, Program Committee, Visions of 2025 and Beyond, ICSE'16, Austin, 2016

Member, Program Committee, 13th IEEE Working Conference on Mining Software Repositories, Austin, 2016.

Member, Program Committee, 10th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Ciudad Real, 2016

Member, Program Committee, 2nd International Workshop on BIG Data Software Engineering, Austin, 2016

Member, Program Committee, 11th International Conference on Global Software Engineering, Orange County, California, 2016.

Member, Program Committee, Foundations of Software Engineering, Bergamo, 2015

Member, Program Committee, 22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering, Osaka, 2015

Member, Program Committee, 31st International Conference on Software Maintenance and Evolution, Bremen, 2015

Member, Program Committee, 9th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Beijing, 2015.

Member, Program Committee, 8th India Software Engineering Conference, Bangalore, 2015

Member, Program Committee, 11th International Conference on Open Source Systems, 2015.

Member, Program Committee, 1st International Workshop on BIG Data Software Engineering, Florence, 2015

Member, Program Committee, 12th IEEE Working Conference on Mining Software Repositories, Florence, 2015.

Member, Program Committee, International Conference on Software Engineering, Hyderabad, India 2014.

Member, Program Committee, 11th IEEE Working Conference on Mining Software Repositories, Hyderabad, 2014.

Member, Program Committee, 8th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Torino, 2014.

Member, Program Committee, Visions and Challenges Track, 22nd ACM SIGSOFT International Symposium on the Foundations of Software Engineering, Hong Kong, 2014

Member, Program Committee, 9th International Conference on Global Software Engineering, Shanghai, China, 2014.

Mentor, Mentoring Program for International Conference on Software Engineering, 2013.

Member, Program Committee, 10th IEEE Working Conference on Mining Software Repositories, San Francisco, 2013.

Member, Program Committee, 7th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Baltimore, 2013.

Member, Program Committee, 8th International Conference on Global Software Engineering, Bari, Italy, 2013.

Member, Program Committee, 9th IEEE Working Conference on Mining Software Repositories, Zürich, 2012.

Member, Program Committee, 6th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Lund, Sweden, 2012.

Member, Program Committee, 7th International Conference on Global Software Engineering, Porto Alegre, Brazil, 2012.

Mentor, Mentoring Program for International Conference on Software Engineering, Zürich, 2012.

Member, Program Committee, 7th International Conference on Global Software Engineering, Porto Alegre, Brazil, 2012.

Member, Program Committee, 8th IEEE Working Conference on Mining Software Repositories, Hawaii, 2011.

Member, Program Committee, 5th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement, Banff, Canada, 2011.

Member, Program Committee, 6th International Conference on Global Software Engineering, Helsinki, Finland, 2011.

Member, Program Committee, 4-th International Symposium on Empirical Software Engineering and Measurement, Bolzano-Bozen, Italy, 2010.

Member, Program Committee, International Conference on Software Engineering, Cape Town, South Africa 2010.

Member, Program Committee, 6th IEEE Working Conference on Mining Software Repositories, Vancouver, Canada, 2009.

Member, Program Committee, 4th International Conference on Global Software Engineering, Limerick, Ireland 2009.

Member, Program Committee, International Workshop on Defects in Large Software Systems, Chicago, USA, 2009.

Expert Review Committee, Canada Foundation for Innovation, 2009.

NSF review panel, 2008.

Member, Program Committee, 5th Working Conference on Mining Software Repositories, Leipzig, Germany, 2008.

Member, Program Committee, Third International Conference on Global Software Engineering, Bangalore, India 2008.

Member, Program Committee, International Conference on Empirical Software Engineering and Measurement, Kaiserslautern, Germany, 2008.

Member, Program Committee, International Workshop on Defects in Large Software Systems, Seattle, USA, 2008.

Member, Program Committee, 4th International Workshop on Mining Software Repositories, Minneapolis, USA, 2007.

Member, Program Committee, International Conference on Empirical Software Engineering and Measurement, Madrid, Spain, 2007.

Member, Program Committee, 3rd International Workshop on Mining Software Repositories, Shanghai, China, 2006.

Member, Program Committee, International Conference on Software Engineering, 2006.

Member, Program Committee, 2nd International Workshop on Mining Software Repositories, St. Lois, Missouri, USA, 2005.

Member, Program Committee, 11th International Symposium on Software Metrics, Como, Italy, 2005.

Member, Program Committee, 6th International Conference on Product Focused Software Process Improvement, Oulu, Finland, 2005.

Member, Program Committee, Third International Workshop on Global Software Development, Edinburgh, Scotland, 2004.

Member, Program Committee, Third International Symposium on Empirical Software Engineering, Redondo Beach, CA, 2004.

Member, Program Committee, Ninth International Symposium on Software Metrics, Sydney, 2004.

Member, Program Committee, Second International Symposium on Empirical Software Engineering, Rome, 2003.

Member, Program Committee, Ninth International Symposium on Software Metrics, Sydney, 2003.

Member, Program Committee, First International Symposium on Symposium on Empirical Software Engineering, Nara, 2002.

Member, Program Committee, Eighth International Symposium on Software Metrics, Ottawa, 2002.

Member, Program Committee, Seventh International Symposium on Software Metrics, London, 2001.

Member, Program Committee, International Conference on Software Maintenance, San Jose, California, 2000.

Member, Program Committee, Sixth International Symposium on Software Metrics, Boca Raton, Florida, 1999.

Member, International Software Engineering Research Network, representative for Avaya.

Senior Member, IEEE, IEEE Computer Society.

Member, American Association for the Advancement of Science.

Software Systems

DRS-OSS: Diff Risk Score prediction tool for OSS projects.

World of Code (WoC) Enabling a research workflow for mining and analyzing the universe of open source version control data.

Risky code prediction: a tool that integrates data from multiple customer support and development systems to model the future chances of a customer-reported defect, identifies one percent of the most risky project code, and suggests the most effective risk remediation methods.

Predicting vulnerable customer systems: a tool that uses customer configuration, service reports, and product issues to model the chances of serious customer dissatisfaction. It presents the list of "hottest" customers with the information necessary to take the necessary action.

Customer quality measure: integrates sales, support, and development data to provide customer perception of software quality: the chances that an average customer will report a valid product issue.

Universal Version History: an index of all versions of the source code for open source projects (presently with over 220M unique file-versions) .

Expertise Browser: a web-based system to assist developers, testers, and managers in identifying experts for a number of software development tasks.

SoftChunk: a system to identify candidate chunks for distributed development across several locations.

Change Risk Predictor: is a tool to model and forecast the likelihood that a software change will cause a failure upon delivery.

SoftChange: a collection of tools to measure and model software changes.

LiveDocs: a framework for easily integrating and controlling information visualization (infoVis) components within web pages to create powerful interactive "live" documents.

FIASCO (Functional Image Analysis Software - Computational Olio): a collection of tools for processing functional Magnetic Resonance Imaging data.

I^3 (Interactive Image Index, pronounced *Ice Cube*): a system to analyze a collection of tens of thousands of images.

MIASMA (Multivariate Interactive Animation System For Map Analysis): a system to model and analyze spread of 57 reportable infectious diseases in the United States.