

Jean-Baptiste Tristan

PERSONAL INFORMATION

Email tristanj@bc.edu
Website <https://jtristan.github.io/>

EDUCATION

Ph.D. computer science, 2009

University of Paris 7, Paris, France

- Title: Formal Verification of Translation Validators
- Performed at INRIA (French Institute for Research in Computer Science and Automation)

M.Sc. computer science, 2006

Ecole Normale Supérieure, Paris, France

Undergraduate studies

I obtained several French diplomas that do not correspond well to US diplomas

- “DEUG” in mathematics and computer science (University of Paris 7)
- “License” in computer science (University of Paris 7)
- “Magistere” in mathematics and computer science (Ecole Normale Supérieure of Paris)

PROFESSIONAL POSITIONS

Boston College , Chestnut Hill, Massachusetts USA <i>Associate Professor</i>	08/2020-present
Harvard University , Cambridge, Massachusetts USA <i>Visiting Lecturer</i>	Fall 2019
Oracle labs , Burlington, Massachusetts USA <i>Consulting Member of Technical Staff</i>	06/2019-07/2020
Oracle labs , Burlington, Massachusetts USA <i>Principal Member of Technical Staff</i>	10/2015-06/2019
Harvard University , Cambridge, Massachusetts USA <i>Visiting Lecturer</i>	Fall 2015
Oracle labs , Burlington, Massachusetts USA <i>Senior Member of Technical Staff</i>	11/2011-10/2015
Harvard University , Cambridge, Massachusetts USA <i>Postdoctoral fellow</i>	11/2009 - 11/2011
Microsoft research-INRIA joint center , Saclay, France <i>Intern</i>	Fall 2009
Harvard University , Cambridge, Massachusetts USA <i>Intern</i>	Summer 2005
Exalead R&D , Paris, France <i>Intern</i>	Summer 2004
University of Paris, 7 , Paris, France <i>Intern</i>	Summer 2003

Awards & Recognition

Co-recipient of the 2021 ACM Software System award.

Keynote speaker at the first international conference on Probabilistic Programming.

Co-recipient of the 2011 La Recherche award in Information Sciences.

Senior member of the ACM.

Invited to the IFIP working group on Functional Programming.

Invited to the IFIP working group on programming languages.

Speaker at the Oracle Product Architect Community.

Panelist at the Oracle Product Architect Community.

Paper Selected for contributed talk at AABI 2021

Spotlight paper at ICML 2021

Spotlight paper at NeurIPS 2014

Spotlight paper at LearningSys 2016

Paper selected for journal publication at PPOPP 2016

Paper selected for journal publication at PPOPP 2017

Grants

Oracle (PI): \$50,000 **2021**
Transfer Learning and Invariance for Bayesian Optimization.

Schiller Institute (co-PI): \$45,500 **2021**
Hierarchical Gaussian Process Regression for Meta-Learning of Molecular Geometry Optimization.

National Science Foundation (PI): \$963,189 **2021**
SHF: Medium: Formally Verified Compilation of Probabilistic Programs.

Oracle (PI): \$50,000 **2020**
Formal Semantics and Verified Parsing for an Inference Language.

Service

UNIVERSITY SERVICE

Member of the Schiller Search Committee (2021-).

Member of the Fitzgerald Search Committee (2021-).

Member of Provost's Advisory Council at Boston College (2021-).

Member of the Boston College Cluster Committee (2020-).

ACADEMIC SERVICE

Steering Committee: POPL workshop on Languages for Inference (LAFI).

Organizer: Colloquium on Probabilistic Programming, Collège de France (2022).

Program Chair: POPL'22 workshop on Languages for Inference (LAFI'22).

General Chair: Third International Conference on Probabilistic Programming (ProbProg'21).

Program Chair: POPL'21 workshop on Languages for Inference (LAFI'21).

Program Chair: Second International Conference on Probabilistic Programming (ProbProg'20).

Program Committee: FMCAD 2021, ASPLOS 2021 ERC member, HOPL 4 PC member, PLDI'18 PC member, PPS'18 PC member, IBM PL day 2016 PC member, SNAPL 2017 PC Member, PAPI 2016 PC Member, PPOPP 2016 PC Member, POPL 2012 External Reviewing Committee, Coq Workshop 2012 PC Member.

Referee: ACM Transactions On Parallel Computing, Communication of the ACM, ACM Transactions On Programming Languages and Systems, ACM Transaction on Architecture and Code Optimization, Software Practice & Experience, Information Processing Letters, Higher-Order and Symbolic Computation.

Reviewer: AISTATS, SOCC, NeurIPS, ICML, POPL, PLDI, PPOPP, DISC, PPDP, SSV, CAV.

Other: National Science Foundation panelist in 2013, 2014, 2015. Treasurer for ICFP 2013.

INDUSTRY SERVICE

Member of Oracle's patent review committee. (2019-2020).

Participated in M&A tech due diligence at Oracle.

Teaching

INSTRUCTOR	Boston College , Chestnut Hill, Massachusetts USA <i>CSCI 1101.02: Introduction to Computer Science</i>	Spring 2022	
	Boston College , Chestnut Hill, Massachusetts USA <i>CSCI 1101.03: Introduction to Computer Science</i>	Spring 2022	
	Boston College , Chestnut Hill, Massachusetts USA <i>CSCI 3340/CHEM 5521: Intro to Machine Learning, Applications to Chemistry</i>	Fall 2021	
	Boston College , Chestnut Hill, Massachusetts USA <i>CSCI 3383: Algorithms</i>	Spring 2021	
	Boston College , Chestnut Hill, Massachusetts USA <i>CSCI 3383: Algorithms</i>	Fall 2020	
	Harvard University , Cambridge, Massachusetts USA <i>CS 281: Advanced Machine Learning</i>	Fall 2019	
	Harvard University , Cambridge, Massachusetts USA <i>CS 153: Compiler Construction</i>	Fall 2015	
	INDEPENDENT STUDIES	Boston College , Chestnut Hill, Massachusetts USA <i>Independent Study: Quantum Chemistry</i>	Spring 2022
		Boston College , Chestnut Hill, Massachusetts USA <i>Independent Study: Applied Natural Language Processing</i>	Spring 2021
Boston College , Chestnut Hill, Massachusetts USA <i>Independent Study: Machine Learning for Chemistry</i>		Fall 2020	
TEACHING ASSISTANT	Harvard University , Cambridge, Massachusetts USA <i>Teaching fellow, CS51: Introduction to computer science II</i>	Spring 2011	
	Harvard University , Cambridge, Massachusetts USA <i>Teaching fellow, CS50: Introduction to computer science I</i>	Fall 2010	

Student Supervision

POSTDOCS	Tarakaram Gollamudi (2022-present).	
	Daniel Huang (2020-2021).	Now assistant professor at San Francisco State University.
PH.D. STUDENTS	Zhen Liang (2022-present) Ph.D. Mathematics, Boston College, co-supervised with Prof. Eli Grigsby.	
	Chong Teng (2021-present) Ph.D. Chemistry, Boston College, with prof. Lucas Bao (primary supervisor).	
	Daniel Huang (2012-2017) Ph.D. Computer Science, Harvard, with prof. Greg Morrisett (primary supervisor).	
	Paul Govereau (2005-2011) Ph.D. Computer Science, Harvard, with prof. Greg Morrisett (primary supervisor).	
RESEARCH ASSOCIATES	Caleb Miller (2021-present).	
	Julian Asilis (2021-2022)	Now Ph.D. student at the University Southern California.
SENIOR THESIS	Jieqi Di (2021-2022). Co-supervised with Kathryn Lindsey, math department. Scholar of the College.	
	Ronan Manvelian (2021-2022).	
	Gina Chun (2021-2022).	
	Darius Russell Kish (2020-2021)	Now a Ph.D. student at Harvard University.
	Bryan Ward (2020-2021)	Now a research engineer at the Flatiron institute.
	Emily Walker (2020-2021)	Now at McKinsey & Company.
INTERNSHIPS	Ananya Barthakur (Boston Collge), Gina Chun (Boston College),Changee Park (KAIST), Joe Tassarotti (Harvard), Daniel Huang (2x) (Harvard), Manzil Zaheer (2x) (CMU), Sam Anzaroot (Umass Amherst), Jay-Yoon Lee (CMU) ,Koundinya Vajha (U. Pittsburgh), Hao Wu (Northeastern), Chanwei Hu (Duke), Aishwaria Kamath (UMass), Rashika Mishra (Ut Dallas).	

Publications

THESIS & JOURNAL *Geometry Meta-Optimization*

PUBLICATIONS Daniel Huang, Lucas Bao, Jean-Baptiste Tristan
In Journal of Chemical Physics.

mad-GP: Automatic Differentiation of Gaussian Processes for Molecules and Materials

Daniel Huang, Chong Teng, Lucas Bao, Jean-Baptiste Tristan
In Journal of Mathematical Chemistry.

Using Butterfly-Patterned Partial Sums to Draw from Discrete Distributions

Guy L. Steele Jr., Jean-Baptiste Tristan
In **TOPC'19**: ACM Transaction on Parallel Computing, 2019.

Adding Approximate Counters

Guy L. Steele Jr., Jean-Baptiste Tristan
In **TOPC'17**: ACM Transaction on Parallel Computing, 2017.

Formal Verification of Translation Validators

Jean-Baptiste Tristan
Ph.D. dissertation

CONFERENCE
PUBLICATIONS

Computable PAC Learning of Continuous Features

Nathanael Ackerman, Julian Asilis, Jieqi Di, Cameron Freer, Jean-Baptiste Tristan
In **LICS'22**: Thirty-Seventh Annual ACM/IEEE Symposium on Logic in Computer Science.

Conjugate Energy-Based Models

Hao Wu, Babak Esmaeili, Michael L Wick, Jean-Baptiste Tristan, Jan-Willem van de Meent
In **ICML'21**: International Conference on Machine Learning, 2021. **Spotlight**

Rate-Regularization and Generalization in Variational Autoencoders

Alican Bozkurt, Babak Esmaeili, Jean-Baptiste Tristan, Dana Brooks, Jennifer Dy, Jan-Willem van de Meent
In **AISTATS'21**: The 24th International Conference on Artificial Intelligence and Statistics, 2021.

A Formal Proof of PAC Learnability for Decision Stumps

Joseph Tassarotti, Koundinya Vajjha, Anindya Banerjee, Jean-Baptiste Tristan
In **CPP'21**: Certified Programs and Proofs, 2021.

Conjugate Energy-Based Models

Hao Wu, Babak Esmaeili, Michael L Wick, Jean-Baptiste Tristan, Jan-Willem van de Meent
In **AABI'21**: 3rd Symposium on Advances in Approximate Bayesian Inference, 2021. **Talk**

Online Post-Processing In Rankings For Fair Utility Maximization

Ananya Gupta, Eric Johnson, Aditya Kumar Roy, Justin Payan, Ari Kobren, Swetasudha Panda, Michael Wick, Jean-Baptiste Tristan.
In **WSDM'21**: The ACM 14th International Conference on Web Search and Data Mining, 2021.
Talk

Unlocking Fairness: a Trade-off Revisited

Michael L. Wick, Swetasudha Panda, Jean-Baptiste Tristan.
In **NeurIPS'19**: 33rd Conference on Neural Information Processing Systems, 2019.

Scaling Hierarchical Coreference with Homomorphic Compression

Michael L. Wick, Swetasudha Panda, Joseph Tassarotti, Jean-Baptiste Tristan.

In **AKBC'19**: 1st Conference on Automated Knowledge Base Construction, 2019.

Sketching for Latent Dirichlet-Categorical Models

Joseph Tassarotti, Jean-Baptiste Tristan, Michael L. Wick.

In **AISTATS'19**: International Conference on Artificial Intelligence and Statistics, 2019.

Gradient-based Inference for Networks with Output Constraints

Jay-Yoon Lee, Sanket Mehta, Michael L. Wick, Jean-Baptiste Tristan, Jaime Carbonell.

In **AAAI'19**: Thirty-Third AAAI Conference on Artificial Intelligence, 2019.

Flexible Compilation of Probabilistic Programs

Daniel Huang, Jean-Baptiste Tristan, Greg Morrisett.

In **PLDI'17**: ACM SIGPLAN Conference on Programming Language Design and Implementation, 2017.

Using Butterfly-Patterned Partial Sums to Optimize GPU Memory Accesses for Drawing from Discrete Distributions

Guy Steele, Jean-Baptiste Tristan.

In **PPOPP'17**: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2017.

Exponential Stochastic Cellular Automata for Massively Parallel Inference

Manzil Zaheer, Michael Wick, Jean-Baptiste Tristan, Alex Smola, Guy Steele.

In **AISTATS'16**: International Conference on Artificial Intelligence and Statistics, 2016.

Adding approximate counters

Guy Steele, Jean-Baptiste Tristan.

In **PPOPP'16**: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2016.

Efficient Training of LDA on a GPU by Mean-for-Mode Estimation

Jean-Baptiste Tristan, Joseph Tassarotti, Guy Steele.

In **ICML'15**: International Conference on Machine Learning, 2015.

Augur: Data-Parallel Probabilistic Modeling

Jean-Baptiste Tristan, Daniel Huang, Joseph Tassarotti, Adam Pockock, Stephen J. Green, Guy Steele.

In **NIPS'14**: Annual Conference on Neural Information Processing Systems, 2014. **Spotlight**

Parallel programming with big operators

Changhee Park, Guy Steele, Jean-Baptiste Tristan.

In **PPOPP'13**: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2013.

RockSalt: Better, Faster, Stronger SFI for the x86

Greg Morrisett, Gang Tan, Joseph Tassarotti, Jean-Baptiste Tristan, Edward Gan.

In **PLDI'12**: ACM SIGPLAN Conference on Programming Language Design and Implementation, 2012.

Evaluating Value-Graph Translation Validation for LLVM

Jean-Baptiste Tristan, Paul Govereau, Greg Morrisett.

In **PLDI'11**: ACM SIGPLAN Conference on Programming Language Design and Implementation,

2011.

A simple, verified validator for software pipelining

Jean-Baptiste Tristan, Xavier Leroy.

In **POPL '10**: ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, 2010.

Verified Validation of Lazy Code Motion

Jean-Baptiste Tristan, Xavier Leroy.

In **PLDI '09**: ACM SIGPLAN Conference on Programming Language Design and Implementation, 2009.

Formal verification of translation validators: A case study on instruction scheduling optimizations

Jean-Baptiste Tristan, Xavier Leroy.

In **POPL '08**: ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, 2008.

WORKSHOP
PUBLICATIONS

Fair Online Post-Processing for Black-Box ML Screening Systems

Swetasudha Panda, Ari Kobren, Jean-Baptiste Tristan, Michael Wick (Oracle Labs)

In **WIML'20**: 15th Women in Machine Learning Workshop.

Using Bayes Factors to Control for Fairness A Case Study on Learning To Rank

Swetasudha Panda, Jean-baptiste Tristan, Haniyeh Mahmoudian, Pallika Kanani, Michael Wick

In **Robust AI in FS'19**: NeurIPS 2019 Workshop on Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy.

Enforcing Output Constraints via SGD: A Step Towards Neural Lagrangian Relaxation

Jay-Yoon Lee, Michael L. Wick, Jean-Baptiste Tristan, Jaime Carbonell

In **AKBC'17**: Workshop on Automated Knowledge Base Construction, 2017.

Sketchy LDA: Towards Streaming Inference

Jean-Baptiste Tristan, Michael L. Wick, Joseph Tassarotti

In **ML Systems'17**: Workshop on ML Systems, 2017.

Comparing Gibbs, EM and SEM for MAP Inference in Mixture Models

Manzil Zaheer, Michael Wick, Satwik Kottur, Jean-Baptiste Tristan.

In **OPT'15**: Optimization for Machine Learning, 2015.

Exponential Stochastic Cellular Automata for Massively Parallel Inference

Manzil Zaheer, Michael Wick, Jean-Baptiste Tristan, Alex Smola, Guy Steele.

In **LearningSys'15**: Workshop on Machine Learning Systems, 2015. **Spotlight**.

Patents

Ensembled decision systems using feature hashing models

Jean-Baptiste Tristan, Adam Pockock, Michael Wick, Guy Steele.

Data-parallel parameter estimation of the Latent Dirichlet allocation model by greedy Gibbs sampling

Jean-Baptiste Tristan, Guy L. Steele Jr.

Systems and Methods for Scalable Hierarchical Coreference

Michael L. Wick, Jean-Baptiste Tristan, Guy L. Steele Jr.

Data-Parallel Probabilistic Inference

Jean-Baptiste Tristan, Guy L. Steele, JR., Daniel E. Huang, Joseph Tassarotti

Learning topics by simulation of a stochastic cellular automaton

Jean-Baptiste Tristan, Stephen J. Green, Guy L. Steele, Jr., Manzil Zaheer

Parallel Gibbs sampler using butterfly-patterned partial sums

Guy L. Steele, Jr., Jean-Baptiste Tristan

Method and system for latent dirichlet allocation computation using approximate counters

Guy L. Steele, Jr., Jean-Baptiste Tristan

Method and system for distributed latent dirichlet allocation computation using addition of approximate counters

Guy L. Steele, Jr., Jean-Baptiste Tristan

Sparse and data-parallel inference method and system for the latent Dirichlet allocation model

Jean-Baptiste Tristan, Joseph Tassarotti, Guy L. Steele Jr.