

# Jean-Baptiste Tristan

---

## PERSONAL INFORMATION

Email jean.baptiste.tristan@gmail.com  
Phone (617) 997-1404  
Citizenship French, permanent resident of the United States

## RESEARCH INTERESTS

- Machine learning
- Compilers, program analysis
- Formal verification
- Parallel and distributed learning systems
- Probabilistic Graphical models, probabilistic programming
- Statistical inference, Markov Chain Monte Carlo algorithms
- Program induction

## 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)

## AWARDS

Recipient of the **2011 La Recherche award in Information Sciences** along with Sandrine Blazy, Zaynah Dargaye, and Xavier Leroy for our work on the CompCert verified C compiler.

Invited to the IFIP working group on Functional Programming and the IFIP working group on programming languages.

## RESEARCH EXPERIENCE

**Oracle labs**, Burlington, Massachusetts USA  
*Principal Member of Technical Staff* **10/2015-present**

**Oracle labs**, Burlington, Massachusetts USA  
*Senior Member of Technical Staff* **11/2011-10/2015**

**Harvard University**, Cambridge, Massachusetts USA  
*Postdoctoral fellow* **11/2009 - 11/2011**

**Microsoft research-INRIA joint center**, Saclay, France  
*Intern* **Fall 2009**

**Harvard University**, Cambridge, Massachusetts USA  
*Intern* **Summer 2005**

**Exalead R&D**, Paris, France  
*Intern* **Summer 2004**

**University of Paris, 7**, Paris, France

*Intern*

**Summer 2003**

TEACHING  
EXPERIENCE

**Harvard University**, Cambridge, Massachusetts USA

*Visiting Lecturer, CS 153: Compiler Construction*

**Fall 2015**

**Harvard University**, Cambridge, Massachusetts USA

*Teaching fellow, CS51: Introduction to computer science II*

**Spring 2011**

**Harvard University**, Cambridge, Massachusetts USA

*Teaching fellow, CS50: Introduction to computer science I*

**Fall 2010**

PEER-REVIEWED  
PUBLICATIONS

*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.

*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**.

*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.

ACADEMIC SERVICE *Program Committee*: IBM PL day 2016, AISTATS 2017, SNAPL 2017 PC Member, NIPS 2016 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, NIPS, ICML, POPL, PLDI, PPOPP, DISC, PPDP, SSV, CAV.

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

PATENTS 10 patents approved by Oracle's patent review committee.

TOOLS I have had to use several combinations of tools for various projects.

- Scalable topic modeling: CUDA, MPI, java, Fork-Join
- Recurrent neural networks: python, numpy, theano, scikit
- Formal verification: Coq proof assistant, OCaml