

# BAILE (PETER) CHEN

[peterbc@mit.edu](mailto:peterbc@mit.edu) | (215) 252-6412 | [peterbaile.github.io](https://peterbaile.github.io)

## EDUCATION

---

MIT CSAIL, Electrical Engineering & Computer Science (Cumulative GPA: 5.0/5.0) Sep 2022

- PhD student in Data Systems Group, co-advised by Prof. Mike Cafarella and Prof. Mike Stonebraker
- Coursework: Natural Language Processing, Advances in Computer Vision

University of Pennsylvania, School of Engineering & Applied Science (Cumulative GPA: 3.90/4.00) Sep 2018 - May 2022

- Summa Cum Laude; Bachelor of Science in Engineering in Networked & Social Systems Engineering (NETS)
- Relevant Coursework: Machine Learning, Database and Information Systems, Big Data Analytics, Internet and Web Systems, Scalable Cloud Computing

## PUBLICATIONS

---

- *Searching data lakes for nested and joined data.* Yi Zhang, **Peter Baile Chen**, Zachary Ives. Under submission.
- *Fault-tolerant and transactional stateful serverless workflows.* Haoran Zhang, Adney Cardoza, **Peter Baile Chen**, Sebastian Angel, Vincent Liu, OSDI 2020.

## TECHNICAL REPORTS

---

- *ReSolution: Incrementally maintaining knowledge graphs.* **Peter Baile Chen**, Zachary Ives.

## SELECTED RESEARCH PROJECTS

---

Log analysis project, MIT Data Systems Group Sep 2022 - Present

Machine logs contain rich information about the underlying applications, systems, networks that generated them. However, given its semi-structured nature and enormous volume, extracting useful information from these logs is a challenging task. This project aims to devise methods to have a better understanding of logs through two folds:

- We try to leverage SOTA NLP language models to generate embeddings for log sequences to infer missing values at different time intervals and then answer counterfactual questions (e.g. what would have happened if workload Y was executed instead of workload X)
- As part of the Database-oriented operating system (DBOS) project, we try to leverage correlations between logs at different levels (e.g. network logs and application logs) within a hierarchy to better identify unusual patterns and behaviors for security purposes

Research Assistant, ReSolution project, University of Pennsylvania Database Group Jun 2021 - Present

Inferencing is a common inefficiency of ML models. When changes occur to training data, the model is retrained and weights are updated, leading to recomputation of inference results. This project aims to tackle the prolonged reinferencing time upon retraining by incrementally recomputing inference results. One of the applications is the incremental maintenance of entity matching in knowledge graphs and associated query results.

- *Independent project*; advised by Professor Zack Ives (Department Chair of Computer Science)
- Devised the algorithm that uses existing inference results as threshold for subsequent inference results to achieve pruning
- Extended generality to various ML models for entity matching (e.g. basic deep learning model, deep learning model with BERT/ LSTM)

Research Assistant, Juneau project, University of Pennsylvania Database Group ([github.com/juneau-project](https://github.com/juneau-project)) Jun 2020 - May 2022

Juneau is a data management system for machine learning and data science projects to help data scientists readily create their own pipelines via reusing existing solutions and data from Kaggle/Github. With much improved scalability, the latest version of Juneau supports searching hierarchical and joined data as well as detecting composite data profiles.

- Joint work with Yi Zhang; advised by Professor Zack Ives (Department Chair of Computer Science)
- Developed the recursive parsing algorithm that normalizes hierarchical data into base relational tables
- Initiated the use of sketching techniques to speed up the process of creating and matching data profiles
- Enhanced scalability through pre-computations, incremental techniques, and in-database implementations

*Research Assistant, Beldi system, University of Pennsylvania Distributed Systems Lab ([github.com/eniac/Beldi](https://github.com/eniac/Beldi))* Jun 2019 - Aug 2020

Beldi is a library and runtime system for writing and composing fault-tolerant and transactional stateful serverless functions that addresses the limited applicability of serverless computing to non-trivial end-to-end stateful applications. It extended the existing log-based fault-tolerant approach with new data structures, transaction protocols, function invocations, and garbage collection.

- Joint work with Haoran Zhang; advised by Professor Sebastian Angel and Vincent Liu
- Built the underlying AWS infrastructure and Python APIs for writing and composing stateful serverless functions
- Implemented variants of two-phase locking, two-phase commits, and lock with intent to enable fault tolerance and serializability of concurrent transactions

*Research Assistant, Physics Department, The University of Hong Kong* Aug 2016 - Dec 2016

- Examined the possibility of semiconductor-metal devices being the next-generation non-volatile memory devices
- Investigated the resistive switching behaviors of copper-doped zinc oxide using different metals as electrodes

## PROFESSIONAL EXPERIENCE

---

*Director of Tech Department, the Daily Pennsylvanian* Jan 2020 - Feb 2021

- Revitalized the DP's digital infrastructure with the upgrade of website framework and introduction of ML techniques
- Pioneered the use of modern web frameworks (React, GatsbyJS, NextJS) to replace out-dated templating languages
- Transformed the Project Pages ([projects.thedp.com](https://projects.thedp.com)) from basic HTML to GatsbyJS which renders faster and more interactive web pages
- Programmed and integrated content recommendation engine into the main site ([thedp.com](https://thedp.com))

*Team Lead & Full-Stack Web Developer, Penn Labs, University of Pennsylvania* Feb 2019 - Dec 2020

- Utilized university authentication, dining and laundry APIs to build *Penn Basics* ([pennbasics.com](https://pennbasics.com))
- Created the university's first food truck API, displaying operation hours, locations, menu, ratings, and reviews

*Full-Stack Web Developer, Masterson Technology Limited* Dec 2018 - Jan 2019

- Designed and built an online college application portal for startup *AppTrack Asia*

*Backend Developer, Perfect World Co., Ltd* Jul 2018 - Aug 2018

- Participated in the development and testing of MMO *LieHuoRuGe*

## TEACHING EXPERIENCE

---

*Instructor, CIS 197: Javascript, University of Pennsylvania* Sep 2020 - May 2022

- Give weekly lectures to classes of 30 students, create assignments, and hold office hours
- Best-rated instructor of the course as of May 2022

*Teaching Assistant, University of Pennsylvania* Sep 2019 - May 2022

- CIS 555, Internet and Web Systems; CIS 320, Introduction to Algorithms (created the first course website [www.seas.upenn.edu/~cis320/](http://www.seas.upenn.edu/~cis320/) which has more than 500 page views per week); NETS 212, Scalable and Cloud Computing (mentored student teams in month-long final project to build simplified Facebook)
- Grade assignments, hold weekly office hours, and update course websites

## AWARDS & ACHIEVEMENTS

---

- *CIS Undergraduate TA Hall of Fame, University of Pennsylvania* May 2022
- *Winner, Hult Prize@Penn 2019* Dec 2019
- *Winner, Penn Wharton Innovation Fund Implementation Award* Nov 2018
- *Meritorious Award, 3rd Annual International Mathematical Modeling Challenge* Jun 2017
- *Silver Medalist, International Sustainable World (Energy, Engineering and Environmental) Project Olympiad* May 2016