

# Jeremy Chen

jeremy.yj.chen@gmail.com · www.jeremy-chen.com · github.com/jeremychen0127 · ca.linkedin.com/in/JeremyYChen

---

## EDUCATION

### University of Waterloo

Candidate of Master of Mathematics in Computer Science, **GPA 4.0** 2018 - 2020

- Research Area: graph database cardinality estimation techniques
- Select Coursework: **Advanced Distributed Systems, Data-Intensive Distributed Computing**

Bachelor of Computer Science, Co-operative Program, **GPA 3.9** 2013 - 2018

- Select Coursework: **Distributed Systems, Database Management**, Machine Learning

## WORK EXPERIENCE

**Snowflake Computing**, Data Services Team, Software Engineering Intern Sept. 2019 - Present

- Participate in designing the posting and request components of the data exchange marketplace to allow customers to facilitate data sharing with data consumers in a **distributed setting**
- Design a **synchronization mechanism** for data consumer requests **across deployment shards**

**University of Waterloo**, Data Systems Group, Research Assistant Sept. 2018 - Present

- Design and implement mechanisms in **Java** to estimate cardinality of acyclic queries to improve query processing performance in a **graph database** ([graphflow.io](http://graphflow.io))
- Enabled execution plan visualization of a graph database by implementing the plan component in **Java**

**Uber ATG**, Autonomous Vehicle Maps, Software Engineering Intern Jan. 2018 - Apr. 2018

- Designed a tracking service and **APIs using Apache Thrift** to provide the progress and release details along the production pipeline of the autonomous vehicle (AV) maps
- Taught myself **GoLang** and implemented the map-production tracking system that scales the production pipeline of the AV maps by a factor of at least 10 and minimizes unnecessary human coordination
- Implemented a middle layer in **Python** to integrate the new tracking system into the existing AV map system

**SAP**, Query Processing Team, Software Developer Jan. 2017 - Apr. 2017

- First-authored the [published paper](#) at EDBT 2018 and filed a [patent](#)
- Improved the query processing performance in SAP HANA **relational DB** by a factor of 5 with a new partition pruning algorithm prototyped in **Python**, which avoids unnecessary data partition scans
- Enabled creation of data statistics objects for relational tables in different storage types by adding **SQL** grammar rules to enhance query optimization

**Miovision Technologies**, Organizational Management Product, Full Stack Developer Apr. 2016 - Aug. 2016

- Launched a real-time organizational management system to external enterprise customers
- Implemented messaging, collaborative editing, and subscription components using **Node.js** and **React**

## PROJECTS

**SocialVR - PennApps XV**: visualization of Facebook life events and photos in Google Cardboard VR ([demo](#))

**ARnance - Winner in Hack Western 3**: rendering AR financial dashboard when scanning company logos ([demo](#))