

# Jennifer Campbell

## CONTACT

778-868-5987

jennscampbell@gmail.com

<https://github.com/jennscampbell>

[www.linkedin.com/in/jennscampbell](http://www.linkedin.com/in/jennscampbell)

## SKILLS

### Languages:

- GoLang, Python, PHP, React, JavaScript, Java
- SQL, MongoDB, MATLAB, R

### Software/Libraries:

- AI/ML & Analytics: OpenSearch, SageMaker, Bedrock, Redshift, Databricks, SPSS, Power BI
- Python Libraries: scikit-learn, XGBoost, LightGBM, TensorFlow, PyTorch, spaCy, HuggingFace, OpenCV, pandas, NumPy
- Compute/Hosting: ECS (EC2, Fargate), AWS Lambda, Docker, Kubernetes
- Storage/DB: AWS S3, AWS RDS, DynamoDB
- Monitoring/Logging: AWS CloudWatch, Microsoft Azure
- CI/CD & DevOps: AWS CodePipeline, CodeBuild, CodeDeploy, CodeCommit, GitHub, Jenkins
- Messaging/Queue: RabbitMQ, SQS
- Project Management: Jira, Asan, Slack

## EDUCATION

### B.Sc. Computer Science

University of British Columbia

### Ph.D. Psychology, Quantitative Science (minor)

University of British Columbia

## PROFILE

Senior Machine Learning Engineer with 5+ years building and deploying production ML systems at scale. Experienced in recommendation systems, real-time inference services, and AWS-based ML infrastructure, with a focus on scalable AI applications and end-to-end model deployment supporting millions of requests per hour.

## WORK EXPERIENCE

### Sr. Machine Learning Engineer, *Minute Media* June 2024 – Present

- Built scalable inference services capable of processing over 6 million requests per hour while maintaining sub-200ms latency.
- Engineered and managed MLOps pipelines using AWS services such as SageMaker, OpenSearch, Bedrock, ECS, and Docker to support production AI deployments and semantic retrieval systems
- Optimized containerized inference pipelines to reduce cost by 20% while also reducing deployment time from 15 minutes to under a minute.
- Implemented semantic retrieval systems using OpenSearch vector indexing to power large-scale content matching and recommendation workflows.
- Led development of recommendation and prediction models used for playlist personalization and ad bidding optimization.
- Translated complex ML systems and technical concepts into actionable insights for business and sales teams.
- Delivered production-grade ML services supporting real-time personalization and advertising optimization across major publisher platforms, including Sports Illustrated, FanNation, Players' Tribune, FanSided, 90Min

### Sr. Machine Learning Engineer, *STN Video* Oct 2022 – June 2024

- Orchestrated the entire development lifecycle for a text and video content-matching app, ensuring successful scoping and production release.
- Maintained sub-200ms execution time while handling 2 million calls per hour.
- Administered AWS services, including SageMaker, ECS, ECR, ElastiCache, and Elasticsearch, ensuring efficient utilization and cost optimization.
- Designed and deployed LLM-powered content matching systems leveraging transformer-based architectures and large-scale multimodal datasets.
- Planned and guided testing processes, error logging, and performance maintenance, ensuring system reliability.
- Presented project findings to stakeholders and actively participated in client calls to showcase achievements and address inquiries.

### Machine Learning Engineer, *STN Video* Oct 2021 – Oct 2022

- Used ML methods on massive data sets (e.g., real-time streaming data)
- Extracted, transformed, loaded data from various sources
- Designed and developed best practices and standards around data engineering
- Created data pipelines to allow for seamless updating of model and docker images to staging and production servers

### Machine Learning Specialist, *UBC* Sept 2018 – Sept 2021

- Wrote and optimized ML algorithms to decode neuroimaging and video data
- Developed ETL processes aimed for efficient processing of large matrices
- Carried out data analysis using Python, MATLAB, and R.
- Implemented online experimental studies using Django, Python, JavaScript, RestAPIs, JSON, HTML and CSS.