
BAIDU
RESEARCH
Bai du 百度

BAIDU
RESEARCH

air-job@baidu.com

BAIDU RESEARCH

Co-located in Beijing, Silicon Valley and Seattle, Baidu Research brings together top talents from around the world to focus on future-looking fundamental research in artificial intelligence.

Big Data Lab

BDL works on cutting-edge research that enables big-data-driven AI technology, platform, and service.

Business Intelligence Lab

BIL primarily focuses on developing effective and efficient data analysis techniques for emerging data-intensive applications.

Computational Biology Lab

CBL is still under development and will focus on the application of different computational techniques to detect problems in molecular biology, genomics and drug discovery.

Cognitive Computing Lab

CCL develops technology to the stage of cognitive computing and expands the AI's boundaries.

Institute of Deep Learning

IDL is committed to conducting the cutting-edge research in various areas such as Natural Language Processing, Computational Biology, and Computer Vision.

Institute for Quantum Computing

IQC aims to be a world-class Quantum Artificial Intelligence research strength, and to continuously integrate relevant quantum technologies into Baidu's core business.

Robotics and Autonomous Driving Lab

RAL focuses on robotics and computer vision research, with the applications to autonomous driving, industry and service robots. We aim to build a multi-disciplinary team that combines expertise in computer vision, machine learning, robotics, graphics simulation and human-robot interaction to study holistic solutions. And we bridge the gap between academic research and industrial applications.

Security Research Lab

SRL aims to provide privacy protected computing and security hardening solutions to AI-powered systems.

Silicon Valley AI Lab

SVAIL is advancing state-of-the-art AI with an emphasis on deep learning, speech, natural language processing and high performance computing.

JOIN US

http://research.baidu.com/Career
Send Your CV to: air-job@baidu.com

FOLLOW US

twitter.com/baiduresearch
linkedin.com/company/baidu-usa

Send Your CV to: air-job@baidu.com

BIG DATA LAB

BRIEF INTRODUCTION

Big data, Deep Learning, and huge computing are shaping up AI and are transforming our society. Big-data-driven decision making and automation are being utilized to solve significant challenges faced by our society in an unprecedented scope. At BDL, we are working on cutting-edge research to better harness big data. We have been leading an initiative on "Open and Inclusive AI", aiming to promote equal access to advanced AI capabilities by all parties through significantly reducing the construction and management cost of AI models.

We have developed and maintained Baidu AutoDL, a suite of software and algorithms to use deep learning to design and train deep learning models, which includes Neural Architecture Search, transfer learning, and interpretability of deep learning.

We also have developed FedCube, a secure data sharing platform for federated learning and cloud-based cooperation and computing. It provides users with comprehensive cloud data and optimal scheduling of computing resources and achieves automated and scalable deployment of workflow.

RESEARCH TOPICS

AutoDL Design and Transfer

To facilitate the uses of deep learning, we have designed efficient and effective automatic algorithms for learning deep models by neural architecture search and improving transfer learning performance on real-world datasets.

InterpretDL

To open the black-box deep learning models and improve their transparency, we have developed an open-source toolkit named InterpretDL with many useful interpretation algorithms that provide explanations of deep learning models'behaviors.

FedCube

Several universities and research institutes have used the FedCube platform to analyze big data collected from the Baidu Maps and Baidu search engines. Interesting scientific discoveries have been reported for analyzing the COVID-19 pandemic in China.

BUSINESS INTELLIGENCE LAB

BRIEF INTRODUCTION

At BIL, we design world-leading AI algorithms and models to support real-world applications in Business Intelligence, Smart City, Computational Biology, Data Mining and Machine Learning, aiming to perform cutting-edge research to advance the state-of-the-art.

BIL provides a transparent, supportive and collaborative environment to boost the project success. We collaborate closely with different teams inside or outside the company, and also have strong ties to academic communities.

We are looking for both software engineers who have rich project experience and research scientists who have a track record of research excellence. Our team consists of elite, researchers and engineers, who devote efforts to tackling the most challenging, meaningful projects in AI and related fields. Welcome to join BIL family!

RESEARCH TOPICS

Smart City

In the field of smart city, we have developed city profile and urban knowledge graph, which have been used for development interpretation, trend prediction and provided analysis for urban planning.

Smart Transportation

In the field of smart transportation, we have released the first enterprise-level multi-modal transportation recommendation engine which is applied in Baidu Maps.

Computational Biology

In the field of Computational Biology, we have developed several novel ML-based computational biology algorithms such as drug property prediction, binding affinity prediction and drug design, aiming to provide an AI-enhanced tool for facilitating drug discovery.

Send Your CV to: air-job@baidu.com

COGNITIVE COMPUTING LAB

BRIEF INTRODUCTION

Cognitive Computing Lab is committed to forward-looking research of cognitive intelligence, through in-depth research on large-scale machine learning and mining, reinforcement learning, knowledge mining, and reasoning. We are pushing artificial intelligence technology to the stage of cognitive computing.

The missions of our lab include:

- Working closely with Baidu's core product teams including search & cloud to improve our products through advanced AI technologies;
- Conducting original basic research in machine learning, knowledge mining and others related to cognitive computing;
- Training the next generation of researchers through our postdoctoral program.

RESEARCH PROJECTS

Machine learning & deep learning theory and algorithms

Our team has a strong workforce and has made achievements in machine learning research, including (1) the rigorous theory for autoencoders via an exact characterization in the limit of large dimensions, which revealed interesting phase transition phenomena and provided quantitative answers and insights to several questions that were yet fully understood in the literature, (2) several new contributions for training energy-based generative models (EBM) along with the best performance among all EBMs, (3) a general framework for highly efficiently training large-scale nonlinear machine learning models. These are published in ICLR'19, ICLR'21, NeurIPS'19, etc.

Efficient search, data mining, and knowledge mining algorithms

We developed a large number of novel algorithms for approximating near neighbor search and maximum inner product search and used them for the company's core products in search. We are among the first groups who developed many novel algorithms for open-information extraction, knowledge representations, etc.

INSTITUTE OF DEEP LEARNING

BRIEF INTRODUCTION

Baidu Institute of Deep Learning (IDL) is committed to conducting the cutting edge of research in various areas such as Natural Language Processing, Computational Biology, Human-centered Computing, Intelligent Short Videos and other fundamental problems and emerging topics and Computer Vision. Our goal is to explore cutting-edge algorithms and technologies that can make impact on the research community and industrial products. Our team has been published in top-tier conferences and journals, and delivered research outcomes to Baidu's commercial products.

RESEARCH PROJECTS

Natural Language Processing

Our interest includes various language generation techniques across different modalities such as machine translation, sequence structural prediction, general speech-related tasks, and multimodal multilingual pretraining techniques.

Computational Biology

Our interest includes structure-centered fundamental research, AI-powered drug discovery and next-generation vaccine study.

Human-centered Computing

In our visual computing, one important area is about visual analysis of humans, such as face recognition, presentation attack detection, human aging analysis, affective computing.

Intelligent Short Videos

Current research topics include semantic video understanding, video summarization, talking face synthesis, and action/behavior recognition.

Fundamental problems and emerging topics

We also focus on fundamental problems and emerging topics in computer vision, such as image/ video segmentation, object detection, self-supervised learning, long-tail computing, and new techniques for human-computer interaction.

Send Your CV to: air-job@baidu.com

INSTITUTE FOR QUANTUM COMPUTING

BRIEF INTRODUCTION

Quantum computing is believed to be the heart of the next-generation computing technology. Our mission at Baidu Research is to be a world-class Quantum Artificial Intelligence (AI) research strength, and to continuously integrate relevant quantum technologies into Baidu's core business.

Our goals include building quantum technology capacity; cultivating top talents; leading grand research projects; establishing thought leadership; developing key applications.

We have general interests in Quantum Information Science with three research priority topics: Quantum AI, Quantum Algorithm, and Quantum Architecture - altogether form the QAAA research plan. We aim to build a sustainable quantum ecosystem, and strive to achieve the vision that "Everyone Can Quantum".

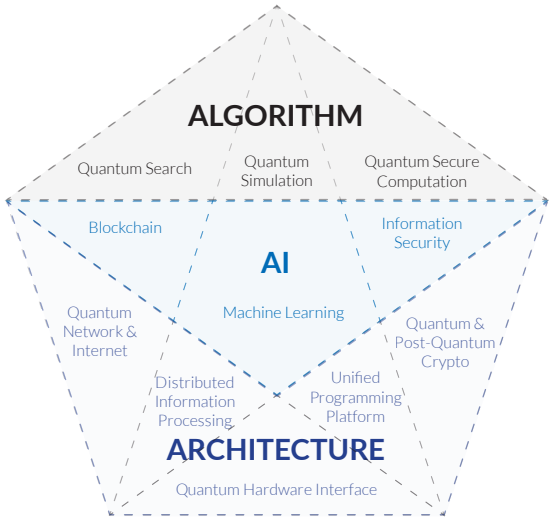

QAAA Research Plan

RESEARCH TOPICS

Quantum AI, Quantum Algorithm, Quantum Architecture.

Quantum AI

Aims: To utilize the advantages of information processing in quantum computing to promote the development of artificial intelligence

To break through the bottleneck of R & D in the quantum area with current artificial intelligence technology

Key Areas: Machine Learning, Information Security, Blockchain

Quantum Algorithm

Aims: To design efficient quantum algorithms for specific tasks

To extend classical algorithm design and analysis techniques into quantum scenario

To optimize existing quantum (classical) algorithms and determine their feasibility and limitations

Key Areas: Quantum Simulation, Quantum Search, Quantum Secure Computation

Quantum Architecture

Aims: To support quantum artificial intelligence and quantum algorithms

To provide a comprehensive Quantum infrastructure as a Service(QaaS)

Key Areas: Quantum Hardware Interface, Distributed Quantum Information Processing, Unified Programming Platform, Quantum Network and Internet, Quantum and Post-Quantum Crypto

ROBOTICS AND AUTONOMOUS DRIVING LAB

BRIEF INTRODUCTION

RAL focus on frontier research and incubation projects, include auto-driving, robotics, perception, navigation and simulation. We aim to bridge the gap between academic research and industrial applications.

RAL team has published research papers, including Science Robotics, TPAMI, TIP, ICCV, CVPR, ECCV, NeurIPS, AAAI, IJRR, ICRA and IROS. The team is also developing novel industry and service robotics systems and applications, including autonomous construction robots.

We have developed novel industry robotics systems and applications, including autonomous excavators. We strive to be the first to bring research lab results, either internally or externally, to market.

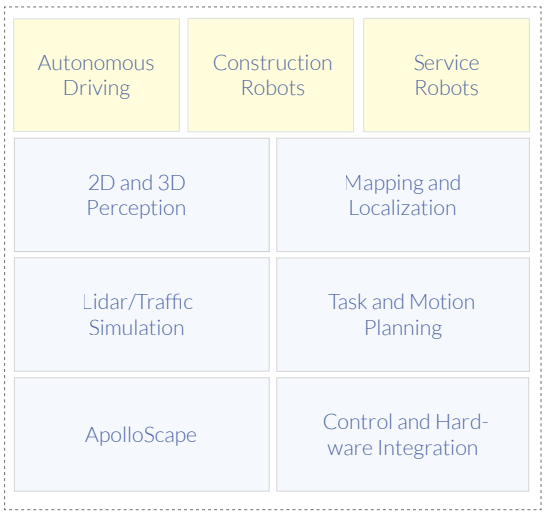

RESEARCH TOPICS

Perception

- Vision for robotics and autonomous vehicles
- Point cloud and depth analysis
- Low-level and physics-based vision
- 3D reconstruction, mapping and localization
- Novel applications

Planning and Control

- Task planning
- Motion planning
- Trajectory optimization
- Robust control

Applications

Autonomous Driving

Autonomous Excavator

Send Your CV to: air-job@baidu.com

SECURITY RESEARCH LAB

BRIEF INTRODUCTION

Cybersecurity and data privacy are important issues for our digitized world which increasingly depends on software, hardware, and cloud infrastructure. In Baidu Security Research Lab, we provide world classic research aiming to address these urgent challenges faced by Internet users, protecting them from cyber-attacks, and guarding user privacy against unauthorized accesses.

Specifically, we conduct public research on topics including designing and building privacy-protected computing solutions, AI robustness evaluation toolbox, and security hardening solutions to protect AI-powered systems such as autonomous driving and smart devices. Our research findings have been published in top-tier security conferences, and we also release cut-edge security capabilities in open-source projects.

RESEARCH TOPICS

Safety and Security of AI Systems

Including AI model robustness assessment and enhancement, security risks assessment of machine learning systems, as well as how to protect AI systems from adversarial ML attacks. Delivering practical protections from AI perspectives in popular scenarios such as autonomous driving.

Privacy Protected Computing

Developing open-source software framework for trusted execution environments, such as SGX, SEV and TrustZone. Ensuring data privacy through a combination of hardware execution environments, verified memory-safe software, and provable attestations.

Software Security Verification

Checking system software implementations against their design specifications, applying formal methods, machine learning, symbolic execution techniques in the area of fuzzing and software verification.

SILICON VALLEY AI LAB

BRIEF INTRODUCTION

SVAIL (Silicon Valley AI Lab) is working on applications in Speech and Language. All of this work intersects with Machine Learning. Much of our recent work has been attempting to understand some recent successes in deep nets and prepare for future successes in months/years/decades to come:

- Language & Machine Learning: Contextual embeddings such as ERNIE and BERT produce excellent results (topping the GLUE leader board). What are these nets doing? Why do they work as well as they do? Can we improve them when we understand better?
- Speech & Machine Learning: Pretrained models like contextual embeddings (ERNIE/BERT) and wave2vec appear to be useful for many tasks in Speech. But is the ultimate solution: "One model to rule them all"? Are there cases helpful in complement models such as wave2vec with knowledge?
- Fundamentals of Task Space and Model Similarities: Both the speech and language projects involve multiple tasks. Can we take advantage of models that are trained or finetuned from different tasks? When will domain transfer be successful?

RESEARCH TOPICS

Speech & Machine Learning

As mentioned above, we are applying pre-trained models to several tasks in Speech. Many of these tasks can be viewed as classification tasks:

1. Classify subjects in an Alzheimer's benchmark as treatment or control
2. Classify utterances in an emotion benchmark as happy, angry, sad, or neutral
3. Classify Chinese syllables by tones: 1, 2, 3, or 4

Other tasks can be viewed as regression problems, such as predicting the duration of words in reading speech. Some pre-trained models, such as ERNIE and BERT, are trained on texts, and other pre-trained models, such as wave2vec, were trained on audio. Both appear to be helpful for these tasks. Pauses are essential for duration modeling because of phrase-final lengthening and Alzheimer's classification because treatment subjects pause while a word is stuck on the tip of their tongue. We have modified ERNIE and BERT to train on text with pauses.

Language & Machine Learning

Most embeddings are based on collocations and Firth's "You shall know a word by the company it keeps." Collocations group words that share semantic classes such as colors (red/yellow) and antonyms (good/bad). For some tasks, such as VAD prediction, it is helpful to distinguish words within groups. We propose a novel embedding to separate red/yellow and good/bad based on translations. Different colors and antonyms are unlikely to share translations in other languages. Therefore, combinations of collocations and translations are better together than either by itself for a VAD prediction task. VAD (valence, arousal & dominance) can be viewed as a generalization of sentiment analysis where good is positive and bad is negative, though VAD uses three dimensions instead of just one.


Bai du 百度