

【职位名称】

AI 大模型算法工程师 / 多模态大模型算法工程师 / 大模型微调 / AI 算法工程师

你好，感谢在超过 80 亿个节点的 Graph 上与我建立边，我是 Tiger，现在带领小伙伴们使用最前沿的开源与闭源大模型以及各类相关技术进行 AI 商业化的创新与创业。

以下是我的一些思考，如果你感兴趣，也欢迎和我深度探讨：

【重要假设】

首先我们假设在不远的未来，每个个体（Customer）都将有自己的 AI 助手或 Agent，我们称之为 CAI；同时每个盈利或非盈利组织（Business）也均会有自己的 AI 助手或 Agent，我们称之为 BAI；在以上两个假设满足的前提下，我们认为，未来 B-C、B-B 甚至 C-C 的交互将主要由代表各自利益的 AI 助手或 Agents 完成。目前已有研究者提出不同的通信协议，从最简单的 MCP 到去除 NLP 层的权重直接交互。这将极大地增加做决策时的信息交流和处理的效率，使其从人类现有的最高 10bit/s 增长三到五个数量级。不难推理出，决策模型将开始从碳基生物的脑中（Neural Network）转移到硅基生物的权重（Artificial Neural Network）中。

【技术角度】

我们重点开发 BAI 用以与 CAI（如 Deepseek / GPT / Claude / Gemini / Qwen / 豆包 / Kimi 等）交互。为了实现这一目标，我们需要进行非结构化信息结构化（KG / GraphRAG / Ontology 等）、人类交互方式 AI 化（MuseTalk / Facefusion / Whisper 等）、AI 人设从静态转为动态化（如星际穿越中的 TARS 可动态且量化的调整幽默度与坦诚程度）、计算量从 Test Time Compute 向训练移动（GCN / GAT / SFT / RLHF / Post-pretraining 等）等多方面的研发与开发。

【商业应用角度】

我们将新突破的 AI 技术应用到之前老技术无法覆盖的场景，从而实现增量市场覆盖，而这部分增量市场体量大约是现存全部电商（Amazon|EBay|淘宝|天猫|京东|PDD|美团|抖音电商等）体量总和的 3 到 9 倍。同时我们认为 BAI 的成功实现会极大的威胁电商或任何中间商的存活，因为 CAI 会跳过需要抽取佣金的中间商而直接与服务或物品提供商的 BAI 对接，从而为其主人——C 端用户提供性价比最高的服务或物品。

【AGI 角度】

我们坚定的相信 AGI 会在我们的有生之年实现，但是我们认为其实现无法由单一个体、单一机构或单一组织通过创造单一 AI 单独实现，只有【多种类垂直 AI 协作】配合【多学科、多组织、多国合作】才能实现【真 . AGI】。而我们第一阶段目标是管理所有的信息（不只是数据），使得信息的产生、合成、运输、收集、存储、处理、交互等均为 AI Native，并等待第二阶段重要 Milestone 的到来（我们认为会诞生一种 AI 原生交互系统 IS - Artificial Intelligent Interactive System 从而取代现有 PC 端和移动端 OS），届时再根据具体情况调整我们的研发

战略与战术。

【目前开放的全职及实习的研究方向为（均为单模态或多模态）】

预训练与微调 Fine-tuning / Post-pretraining with reasoning

垂直知识领域对话机器人 Domain Specific Chatbot

知识图谱构建 Knowledge Graph Construction without Complete Ontology

文档检索 Document Retrieval

基于图的检索增强生成 RAG Graph-based RAG

图神经网络相关 GNN / GCN / GraphSAGE / PinSAGE / GAT

多智能体协作系统 Multi-agent System (MAS)

AI 原生前端与后端 AI Native Frontend / Backend

视频智能讲解或混剪 Video Narration Generation and Auto Editing

AI 虚拟数字人唇纹、表情与动作 AI Avatar

营销全流程自动化 Ai-driven Marketing Automation

其他子方向持续增加中...

【岗位职责】

1. 负责内部新项目 **prototype** 的研发；
2. 阅读最新的论文，学习最前沿的技术/算法/基础架构；
3. 解决别人解决不了的问题。

【岗位要求】

1. 人工智能 或 计算机科学 或 软件工程类；
2. 在校成绩前 10%及以上；
3. 熟练掌握 Python；
4. 熟悉至少一款主流的开源或闭源大模型；
5. 热爱人工智能，不惧怕学习新技术。

【你能获得】

1. 全权负责子项目：你做的不是"某个模块的第 37 版迭代"，而是真正能写进简历的独立项目，你不再只是参数调试的“螺丝钉”，而是从需求定义、技术选型到实验验证的项目主人翁，拥有完整决策与执行权限，让你的想法快速变为成果；
2. 扁平直带：团队无上下级隔阂，技术大牛直带，5 分钟内即可碰撞想法、调整方向，免去漫长审批与沟通干扰；
3. 高效专注 & 真正 Work-Life Balance：我们坚决拒绝无意义加班，八小时深度工作，下班后彻底离线，崇尚“专注 > 内卷 & 效率 > 时长”，让你既能高效产出，也能全身心享受生活；
4. 充足算力保障：只要你需要，我们提供足够的 GPU/TPU 计算资源，训练、微调、推理全流程无卡顿，让你的大模型探索与实验零束缚；
5. 新模型优先体验：任何一款最新发布或 Preview 状态的大模型，你都能第一时间在公司

平台上试用、测试与改进，持续站在 AI 前沿。

【实习岗位薪资】

在读本科实习生 - 人民币 400 – 600 / 天

在读硕士实习生 - 人民币 400 – 700 / 天

在读博士实习生 - 人民币 500 – 1000 / 天

【工作地址】

中国广东省深圳市深圳湾科技生态园

【投递方式】

联系人: Whitney

投递邮箱: talentattraction@linksome.com

【One More Thing】

我们是国内极少数的进行【AI 全方面商业化落地】且同时【坚持长期主义潜心做研发】的商业机构，如果你想要了解更多，发个私信和我聊一聊吧~

【以下的词能帮助更多的小伙伴找到我们】

AI, 人工智能, Artificial Intelligence, 机器学习, Machine Learning, Neural Network, 神经网络, 深度学习, Deep Learning, Transformer, Data Mining, 数据挖掘, 自注意力, Self-Attention, 注意力机制, CNN, 卷积神经网络, RNN, Recurrent Neural Network, 循环神经网络, LSTM, BERT, GPT, Diffusion Model, 扩散模型, 生成式 AI, Generative AI, 大语言模型, Large Language Model, LLM, 图神经网络, Graph Neural Network, GNN, GCN, 图卷积网络, GAT, 图注意力网络, Graph, Knowledge Graph, 知识图谱, Ontology, Vector DB, Vector Database, 向量数据库, Pre-training, 预训练, Post-training, 增量预训练, Continue Pre-training, RAG, Retrieval-Augmented Generation, Embedding, 嵌入表示, Rerank, Computer Vision, 计算机视觉, NLP, 自然语言处理, ASR, 语音识别, TTS, 语音合成, 强化学习, Reinforcement Learning, RLHF, 人类反馈强化学习, Federated Learning, 联邦学习, Edge AI, PyTorch, TensorFlow, CUDA, GPU, Tensor, 张量, Safetensors, Multi-agent, 多智能体, Agent, 代理, Prompt Engineering, 提示工程, Chain of Thought, 思维链, Zero-shot, One-shot, Few-shot, In-Context Learning, Self-supervised Learning, 自监督学习, Unsupervised Learning, 无监督学习, Semi-supervised Learning, 半监督学习, Contrastive Learning, 对比学习, Vision Transformer, 生成对抗网络, GAN, Variational Autoencoder, 变分自动编码器, Autoencoder, 自动编码器, Data Augmentation, 数据增强, Knowledge Distillation, 知识蒸馏, Transfer Learning, 迁移学习, Active Learning, 主动学习, LoRA, 参数高效微调, PEFT, Fine-tuning, 微调, Prompt Tuning, 提示微调, Stable Diffusion, Edge Computing, 边缘计算, 因果推断, Robotics, Tokenization, Hyperparameter, 超参数, Overfitting, 过拟合, Hallucination, 幻觉, Multimodal, 多模态, YOLO, SVM, RPA, Hugging Face, LangChain,

LlamaIndex, Kubernetes, AGI, Data, 数据, Benchmark, Computer Science, 计算机科学, Data Science, 数据科学, Pattern Recognition, 模式识别, Classification, 分类, Regression, 回归, Clustering, 聚类, Dimension Reduction, 降维, Loss Function, 损失函数, Activation Function, 激活函数, Gradient Descent, 梯度下降, Backpropagation, 反向传播, TPU, CPU, 显存, Causal Inference, Computer Vision, 计算机视觉, Feature Engineering, 特征工程, Algorithm, 算法, Sentiment Analysis, 情感分析, GraphSAGE, PinSAGE, Keras, Scikit-learn, OpenCV, Spark MLlib, Docker, Distributed Computing, 分布式计算, Segmentation, 分割, KV Cache

英文版

Position Title

AI Research & Engineering Intern (AI 实习生)

Job Description

From Tiger, Founder @ Linksome

Within a graph of eight-billion people stretching across space-time, the fact that our two nodes formed an edge is nothing short of amazing. I'm grateful for the statistically miraculous coincidence that brought you here.

I'm Tiger, and I currently lead a passionate team focused on AI commercialization — from foundational algorithm exploration to practical business deployment. We work with the latest open-source and closed-source models and build real-world, production-grade applications around them.

Let me share a few ideas we're working on. If any of this resonates with you, send me a message and let's talk.

Core Assumptions

We believe that in the not-so-distant future:

- Every individual (Customer) will have a personal AI agent — we call it CAI.
- Every organization, for-profit or not (Business), will also have its own AI agent — we call it BAI.

Given this, we predict that most B2C, B2B, and even C2C interactions will be handled through AI-to-AI communication, each acting on behalf of their human stakeholders.

Some researchers have already proposed AI communication protocols — from the popular but simple MCP to weight-level interactions that remove the layers for NLP altogether. This could increase information exchange and process efficiency by 3–5 orders of magnitude beyond the human cognitive cap (~10 bits/s). As a result, decision-making will TRULY shift from the carbon-based human brains (neural networks) to silicon-based AI weights (artificial neural networks), for the very first time.

Technical Vision

We're building BAI systems that can naturally and intelligently interact with mainstream CAIs (e.g., GPT, Claude, Gemini, DeepSeek, Qwen, etc.).

This requires solving challenges across multiple fronts:

- Structuring unstructured data (e.g., KG, GraphRAG, Ontology)
- Making human-like communication AI-native (e.g., MuseTalk, Whisper, FaceFusion)
- Designing dynamic AI personas, like *TARS* from *Interstellar*, capable of adjusting attributes like humor or honesty on the fly
- Shifting computational load from test-time to train-time, leveraging methods like GCN, GAT, SFT, RLHF, and post-pretraining strategies

Commercial Thesis

We apply new AI breakthroughs to domains previously unreachable by old-school AI — unlocking incremental markets we estimate to be 3–9 times the total volume of today's global e-commerce (Amazon, Taobao, JD, TikTok Commerce, etc.).

In the long run, we believe BAI agents will disintermediate traditional e-commerce platforms. Why pay a commission to a middleman when your CAI can directly negotiate with the BAI of a product/service provider?

AGI Outlook

We firmly believe AGI will be realized in our lifetime, but not by any single organization or one supermodel.

True AGI will only emerge through a community of AIs just like us:

- Collaboration among many specialized vertical AIs
- Multidisciplinary, multi-organization, international cooperation

Our Phase 1 mission is to make information itself AI-native — from generation to interaction. Phase 2 will begin with the rise of a new form of Operating System: the AI-native Interactive System (IS) — replacing both desktop and mobile interfaces (the old OS) as we know them today. We will adjust our directions accordingly then.

Current Available Research Areas (Single-Modal or Multi-Modal):

- Fine-tuning / Post-pretraining with Reasoning
- Domain-specific Chatbots
- Knowledge Graph Construction w/o Complete Ontology
- Document Retrieval
- Graph-based RAG
- GNN / GCN / GraphSAGE / GAT
- Multi-Agent Systems (MAS)
- AI-native Frontend / Backend
- Video Narration Generation & Auto Editing
- AI Avatar
- AI-driven Marketing Automation

... with new directions constantly emerging

Internship Role: AI Research & Engineering Intern

What You'll Do

Lead prototyping for new internal AI projects (within the areas mentioned above).

Stay up-to-date with cutting-edge techs, and transform novel ideas into productivity systems.

Tackle problems that most people think cannot be solved.

What We're Looking For

Major in AI, Computer Science, Software Engineering or highly related fields.

In the top 10% at university.

Proficiency in Python.

Familiarity with at least one mainstream foundation model (open-source or closed-source).

Deep passion for AI with strong self-motivation to learn new technologies.

What You'll Gain

Real Ownership – Not "Version 37 of a minor module." You'll build something from scratch, own decisions end to end, and drive it into reality.

Flat Structure, Direct Mentorship – No long chains of approval. You'll work directly with technical leaders, iterate fast, and move faster.

Focused Work, No Burnout – We reject fake hustle. Work smart for 8 hours, then fully disconnect. Focus beats grind.

Ample Compute – You'll get the GPU/TPU resources you need, when you need them — for training, finetuning, inference, and everything in between.

Early Access to New Models – Whenever a new foundation model drops, you'll be one of the first to use it, test it, and make it better.

Compensation

Undergraduate students: RMB 400–600/day

Postgraduate students (Master): RMB 400–700/day

Postgraduate students (PhD): RMB 500–1000/day

Work Location

Shenzhen Bay Technology Ecological Park, Shenzhen, Guangdong Province, China

How to Apply

Contact: Whitney

Email: talentattraction@linksome.com

One More Thing

Per Aspera Ad Astra. We're one of the very rare startups in China equally committed to full-stack AI commercialization and long-termism research.

If you're curious, hungry, and ready to build — we'd love to meet you.

The following keywords can help more people find us

AI, Artificial Intelligence, Machine Learning, Neural Network, Deep Learning, Transformer, Data Mining, Self-Attention, CNN, RNN, Recurrent Neural Network, LSTM, BERT, GPT, Diffusion Model, Generative AI, Large Language Model, LLM, Graph Neural Network, GNN, GCN, GAT, Graph, Knowledge Graph, Ontology, Vector DB, Vector Database, Pre-training, Post-training, Continue Pre-training, RAG, Retrieval-Augmented Generation, Embedding, Rerank, Computer Vision, NLP, ASR, TTS, Reinforcement Learning, RLHF, Federated Learning, Edge AI, PyTorch, TensorFlow, CUDA, GPU, Tensor, Safetensors, Multi-agent, Agent, Prompt Engineering, Chain of Thought, Zero-shot, One-shot, Few-shot, In-Context Learning, Self-supervised Learning, Unsupervised Learning, Semi-supervised Learning, Contrastive Learning, Vision Transformer, GAN, Variational Autoencoder, Autoencoder, Data Augmentation, Knowledge Distillation, Transfer Learning, Active Learning, LoRA, PEFT, Fine-tuning, Prompt Tuning, Stable Diffusion, Edge Computing, Robotics, Tokenization, Hyperparameter, Overfitting, Hallucination, Multimodal, YOLO, SVM, RPA, Hugging Face, LangChain, LlamaIndex, Kubernetes, AGI, Data, Benchmark, Computer Science, Data Science, Pattern Recognition, Classification, Regression, Clustering, Dimension Reduction, Loss Function, Activation Function, Gradient Descent, Backpropagation, TPU, CPU, Causal Inference, Computer Vision, Feature Engineering, Algorithm, Sentiment Analysis, GraphSAGE, PinSAGE, Keras, Scikit-learn, OpenCV, Spark MLlib, Docker, Distributed Computing, Segmentation, KV Cache