


You can just do things.

Alex Chambers-Jones

alexcj.co.uk

Survivorship bias.

- Roughly **50%** of startups fail.
- Some of my attempts failed too.
- Had they all failed, you wouldn't be hearing me talking about them!
- AI is crazy, and the world is changing quickly.



I did a lot of programming.

Then I left Berkeley.

- I had a place at Bristol University to study Physics.
- I delayed it by a year for medical reasons — I needed to rest.
- I got bored very quickly.

“That’d be fun.”

- I was interested in languages.
- I decided to make a tool that used fancy probabilistic data structures to estimate the difficulty of sentences on webpages, then translate ones at your level to another language, so you learn by immersion.
- For the first time, I would actually sell it.
- Near-zero strategic thinking.

Metaculus AI Forecasting Tournament

- Metaculus is a prediction platform.
- They were running a tournament to build AI bots to predict the future.
- I didn't do it: "Markets are efficient; I have no edge; it's not worth my time."
- **This was very stupid, because markets are not efficient.**



I built the language tool.

- It was technically good.
- It made some money, but not loads: 20 paying customers.
- My personality is very, very bad at selling to normal people.
- “Don’t sell to consumers” was common wisdom. I knew that and ignored it, and the common wisdom was correct!

But I did some research.

- I needed research, couldn't find it, and did it myself.
- I wrote a blog about what I found.
- It was ridiculously low-hanging fruit.

AboutBlogLanguagesPricingTranslatorLoginGet Started

The more LLMs think, the worse they translate

Deep Translate

Translation that doesn't sound like it. Built on our research.

[Try \\$10 free](#)

Why test this?

In my testing of LLM translation, I've long observed that LLMs tend to perform worse when they think before translating, rather than simply generating the answer directly. I wanted to know if critiquing after translating, rather than thinking before, would act differently, and compare them with a baseline and the technique used in my open-source [hybrid translator](#).

My hypothesis

I previously had some success with my [hybrid translator](#), which performs better than any single model by producing translations with the top 3-5 models, then having GPT-4.1 critique and combine them into a single, very-high-quality translation. I presumed that the critique stage was a major contributor to that success. I thought that applying the same technique to a single LLM output - perhaps with the added variety of two different LLMs - might produce similar improvements. Perhaps diversity of LLMs - having one model critique another - would cover blind spots inherent in the models, and improve performance?

Unexpected results

Company	Thinking	Mean	Std Dev	P90	P99
Aggregate	Ensemble (No Gemma)	89.46	13.93	75.00	28.50
Aggregate	Ensemble (Full)	89.17	15.09	75.00	17.10
Aggregate	Baseline	87.32	16.13	70.00	25.00
Aggregate	Pre-think	83.72	20.84	50.00	10.00
Aggregate	Post-think	83.08	20.57	50.00	10.00

The world is smaller than you think.

- Nobody else was doing AI translation research.
- I was invited to do a paid talk at a company, and random Japanese professors emailed me to chat about it.

Let's try Metaculus.

- I built something super simple.
- It turns out that nobody else was trying.
- I came second place. Good to know! Let's keep at this.

I tried building something to trade Polymarket.

- That failed.
- Sometimes things fail!

A fancy translation system.

- I got a few business customers.
- Then I realised that building a translation company when I don't know any other languages is patently stupid.
- I focused on forecasting instead.

It does translate poetry very well, though. Still a good party trick when I meet someone who actually speaks the language!

Frequency with which each method is preferred in pairwise comparisons.

The screenshot shows a web-based translation interface. At the top, there are two dropdown menus for selecting the source and target languages. The first dropdown is set to 'English' with a small flag icon, and the second is set to 'German' with a small flag icon. Between them is a swap button. Below the language selection is a large text area containing a poem in English. The poem is: 'Tho' much is taken, much abides; and tho' We are not now that strength which in old days Moved earth and heaven, that which we are, we are; One equal temper of heroic hearts, Made weak by time and fate, but strong in will To strive, to seek, to find, and not to yield!'. To the right of the text area is a character count '268/3600'. Below the text area are two buttons: 'Fast' and 'Deep', with 'Deep' being the selected option. To the right of these buttons is a link for 'Advanced Options'. At the bottom of the interface is a large green button labeled 'Translate'.

English ↕ German

Tho' much is taken, much abides; and tho'
We are not now that strength which in old days
Moved earth and heaven, that which we are, we are;
One equal temper of heroic hearts,
Made weak by time and fate, but strong in will
To strive, to seek, to find, and not to yield!

268/3600

Fast Deep Advanced Options

Translate

University was Basically Fine

- I emailed Exa for higher rate limits for my bot.
- The website said higher limits were for business users only. They raised the limits anyway. **Cold emailing works!**
- They offered me an internship, pending legal being OK with a British citizen.
- Then they never got back to me.

Again: failures are very normal.

A random Twitter post.



- Invisible College, by Works in Progress.
- I read enough WIP to write a good application.
- It took 20 minutes. I got in.
- I was way too modest. American norms are different!



Jasmine and Jeremy



Jasmine · 20

- An incredibly accomplished, intelligent AI safety researcher.
- She encouraged me to get into AI safety, explained the routes, and gave me a very useful reference.
- She knows everyone in AI safety.

Jeremy · 20

- An economics student in America.
- He knows everything about AI economics.
- He seems to know everyone in the San Francisco Bay Area.

Lots of other people:


- Isha — AI biosecurity
- Rose — aerospace engineering
- Jack — Cambridge Natural Sciences
- A Swift Institute forecaster
- A former Cabinet Secretary
- Head of robotics at ARIA
- Anthropic's International Affairs lead
- A Jane Street quant
- A Thinking Machines AI researcher
- The Works in Progress team

A week later: San Francisco.

- Jasmine told me about BlueDot Context Week. I applied and went.
- I met lots of really cool people!
- Cultural difference: networking is not a dirty word, and everyone wants to help you succeed.
- I started the extensive application process for MATS, an AI safety research fellowship.

Access + eloquence

- Jasmine invited me to an exclusive event in San Francisco.
- I chatted with senior AI researchers.



“You’ve clearly thought this through a lot.”

“Once we start doing internships, you’re first on my list.”

You Can Just Talk to **Americans**


- Everyone wants to help.
- You can make friends with AI researchers while waiting in line at the airport. I did.
- It's a very different culture. A week in the Bay Area will get you very far.

I said hi to London's AI safety researchers.

- I asked Jasmine for contact details.
- I emailed or texted and asked nicely.
- I went around the MATS and LISA offices.
- It was totally informal. **You can just do things.**
- This probably would've worked without Jasmine too; she just made things smoother.

An opportunity via WhatsApp.

- Jeremy went to a rationalist summer camp in 2025.
- Someone posted in the group chat that they were looking for forecasters for a hedge fund.



I'm not a forecaster.

- I DMed him anyway.
- I pitched that AI forecasting was going to replace human forecasting.
- They should buy my software.

That somehow worked.

- A month of negotiations and presentations.
- We agreed on an exceptionally well-paid internship.
- I'm there until January, then either heading off for MATS or staying in a full-time role.
- I also went to an EA conference and met friends from Cambridge, plus lots of other cool people!

I'm smart, but I'm not a genius.

- I'm a mediocre physics student.
- I am bad at maths relative to my peers, but good at writing.
- I wouldn't stand a chance of getting a normal finance internship.
- Plenty of people in this room are as smart as me.
You can do things too.

UNIVERSITY PHYSICS GRADES

C C C

STRAIGHT CS.

What can we take from this?

YOU CAN JUST DO THINGS.



roon 
@tszzl



say it with me now. experts are fake, smart generalists rule the world,
everything is designed by people no smarter than you, and courage is in
shorter supply than genius

Compete as little as possible.

- Everyone is fighting over the same tiny sliver of possibilities.
- The normal track means thousands of applications and stupid psychological tests.
- Basically everything else has far less competition.

Give it a go.

Cold Emails are Unreasonably Effective

- People will usually reply if you do them well, and they don't take long!
- **Never use AI.**
- Reduce fluff and increase signal, then **KEEP DOING THAT** until it's comically compact.

A BAD COLD EMAIL

Hi John,

I hope this email finds you well. My name is Alex Chambers-Jones, and I'm a Physics student at the University of Bristol. Alongside my degree, I've been pursuing independent programming projects, initially in language learning and translation. I built a browser extension, conducted research into translation quality, and developed a translation system for business customers.

More recently, I've focused on AI forecasting: using large language models to research questions about geopolitical and economic events, and assign probabilities to possible outcomes. This involves combining information retrieval with different models and forecasting approaches. My own system has achieved second place in the Metaculus AI Forecasting tournament.

Just last month, another forecaster, Laertes, beat all the humans in the Metaculus Cup. ManticAI came second, and FutureSearch came fifth. There were 885 forecasters, making 88,558 forecasts on 58 questions. On the current trend, this is going to seriously affect finance: better forecasts could help organisations anticipate economic developments and make better-informed decisions.

Given these developments, and the experience I've gained through my research and commercial projects, I'd be interested in exploring whether there might be an opportunity to discuss an internship or an IP transfer arrangement that could help your organisation get ahead of the competition.

Thank you for your time and consideration.

Best regards,

Alex Chambers-Jones

A GOOD COLD EMAIL

Hi John,

AI recently beat humans at predicting the future, with clear implications for finance. I'm a Physics student at Bristol University, I research AI forecasting, and I've come 2nd place in the Metaculus AI Forecasting tournament. I'd love to discuss an IP transfer or internship to get you ahead of the competition.

Thanks,

Alex Chambers-Jones

Legibility

- A tournament position.
- A blog.
- A project people can actually look at.
- Internships, fellowships, and prior work.
- Even a Twitter account where you're smart.

Put yourself out there in a way that fits in an application.

Eloquence is important. This is gameable.

- Nobody is asking you to do an equation at a party.
- Read a lot. Write a lot. Think about things.
- This isn't about purple prose or English essays.
- Read Paul Graham's essays, Dwarkesh podcasts, and LessWrong.
- Actively increase your **active vocabulary**: the words you draw upon in conversation.

Practice speaking to people.

- Presenting in a boardroom.
- Talking to a powerful person at a party.
- Discussing things in a small group.
- These come up more than talking to a big crowd.

Depending on how I'm doing right now, this may sound hypocritical.

The audience matters.

The success of the “weirdo 20yo” depends heavily on who’s listening.

Industry > Academia

Bay Area > America > Britain > Europe

Tech > Finance > Literally anything else

Startup > Bigco

New > Old

Academia is a particularly hideous status game. Avoid it.

Twitter democratises access.

- In certain niches, it's incredibly useful.
- If you have smart things to say, you will quickly find yourself randomly chatting with very powerful people.
- You find interesting people, research, and opportunities.
- It gives you lots of material and context to discuss.
- The format might not fit you. I'm terrible at posting!

Surround yourself with good people.

- People who are smart, driven, and encouraging.
- People who actually do things and help each other.
- I mostly found them through Twitter, Invisible College, and the Bay Area.

Effective Altruism/Rationalism

- These are big subcultures in the Bay Area, startups, and AI.
- They're very friendly to people inside them.
- Disagreeing with half of EA/rationalism is pretty much what defines you as an EA/rationalist. You don't need to believe it all to benefit from ingroup bias.
- The way you write reveals whether you're an EA/rationalist, so you can silently express it in an email or job application.


What about AI?

- Everything is changing.
- Be willing to change what you're doing. Pivot fast.
- Use AI to learn and build. Keep your own judgement.

[Watch this lecture.](#)

Miles Cranmer / Cambridge
How to Avoid AI Brain Rot

It pulls you to the population mean



There are loads of things to apply to.

- AI safety has a whole pipeline: courses, research programmes, and fellowships.
- Ask your AI to find things that fit your age, interests, and experience.
- Apply to Invisible College too. Your AI might miss it.

BlueDot · SPAR · MATS · ERA · LASR · ARENA · Invisible College

Todo lists and notes you actually USE.

- Keep track of applications, deadlines, and people to reply to.
- Write down ideas you'll otherwise forget.
- Have something you actually use to keep on top of things.

80/20 your entire life

- You can generally get 80% of the way with 20% of the effort.
- Do many things, slightly worse.

But rest when you can

- Preserve the capacity to go 100% for weeks or months when you need to.

Some useful links.

Invisible College / Works in Progress

worksinprogress.co/invisiblecollege

BlueDot / AI courses & programmes

bluedot.org

SPAR / AI safety research

sparai.org

MATS / Research fellowship

matsprogram.org

ERA / Research fellowships

erafellowship.org

LASR / London AI Safety Research

lasrlabs.org

ARENA / Technical AI safety training

arena.education

EAG / Effective Altruism Global

effectivealtruism.org/ea-global

80,000 Hours / Careers

80000hours.org

Paul Graham / What You'll Wish You'd Known

paulgraham.com/hs.html

Miles Cranmer / How to Avoid AI Brain Rot

youtu.be/uXC70uCas2Q

Forecasting questions & competitions

metaculus.com

alexcj.co.uk alexcj2802@gmail.com

Slides: careerspresentation.alexcj.co.uk. I'll ask for them to be posted on Teams.