

Leopold Aschenbrenner, June 2024, *Situational Awareness*

Contents

Introduction 3

History is live in San Francisco.

I. From GPT-4 to AGI: Counting the OOMs 7

AGI by 2027 is strikingly plausible. GPT-2 to GPT-4 took us from ~preschooler to ~smart high-schooler abilities in 4 years. Tracing trendlines in compute (~0.5 orders of magnitude or OOMs/year), algorithmic efficiencies (~0.5 OOMs/year), and “unhobbling” gains (from chatbot to agent), we should expect another preschooler-to-high-schooler-sized qualitative jump by 2027.

II. From AGI to Superintelligence: the Intelligence Explosion 46

AI progress won't stop at human-level. Hundreds of millions of AGIs could automate AI research, compressing a decade of algorithmic progress (5+ OOMs) into 1 year. We would rapidly go from human-level to vastly superhuman AI systems. The power—and the peril—of superintelligence would be dramatic.

III. The Challenges 74

IIIa. Racing to the Trillion-Dollar Cluster 75

The most extraordinary techno-capital acceleration has been set in motion. As AI revenue grows rapidly, many trillions of dollars will go into GPU, datacenter, and power buildout before the end of the decade. The industrial mobilization, including growing US electricity production by 10s of percent, will be intense.

IIIb. Lock Down the Labs: Security for AGI 89

The nation's leading AI labs treat security as an afterthought. Currently, they're basically handing the key secrets for AGI to the CCP on a silver platter. Securing the AGI secrets and weights against the state-actor threat will be an immense effort, and we're not on track.

IIIc. Superalignment 105

Reliably controlling AI systems much smarter than we are is an unsolved technical problem. And while it is a solvable problem, things could very easily go off the rails during a rapid intelligence explosion. Managing this will be extremely tense; failure could easily be catastrophic.

IIIId. The Free World Must Prevail 126

Superintelligence will give a decisive economic and military advantage. China isn't at all out of the game yet. In the race to AGI, the free world's very survival will be at stake. Can we maintain our preeminence over the authoritarian powers? And will we manage to avoid self-destruction along the way?

IV. The Project 141

As the race to AGI intensifies, the national security state will get involved. The USG will wake from its slumber, and by 27/28 we'll get some form of government AGI project. No startup can handle superintelligence. Somewhere in a SCIF, the endgame will be on.

V. Parting Thoughts 156

What if we're right?

Appendix 162

Leopold Aschenbrenner, June 2024

IV. *The Project*

As the race to AGI intensifies, the national security state will get involved. The USG will wake from its slumber, and by 27/28 we'll get some form of government AGI project. No startup can handle superintelligence. Somewhere in a SCIF, the endgame will be on.

U.S. Government

Sensitive Compartmented
Information Facility

We must be curious to learn how such a set of objects—hundreds of power plants, thousands of bombs, tens of thousands of people massed in national establishments—can be traced back to a few people sitting at laboratory benches discussing the peculiar behavior of one type of atom.

SPENCER R. WEART

MANY PLANS FOR “AI governance” are put forth these days, from licensing frontier AI systems to safety standards to a public cloud with a few hundred million in compute for academics. These seem well-intentioned—but to me, it seems like they are making a category error.

I find it an insane proposition that the US government will let a random SF startup develop superintelligence. Imagine if we had developed atomic bombs by letting Uber just improvise.

Superintelligence—AI systems *much* smarter than humans—will have vast power, from developing novel weaponry to driving an explosion in economic growth. Superintelligence will be

the locus of international competition; a lead of months potentially decisive in military conflict.

It is a delusion of those who have unconsciously internalized our brief respite from history that this will not summon more primordial forces. Like many scientists before us, the great minds of San Francisco hope that they can control the destiny of the demon they are birthing. Right now, they still can; for they are among the few with situational awareness, who understand what they are building. But in the next few years, the world will wake up. So too will the national security state. History will make a triumphant return.

As in many times before—Covid, WWII—it will seem as though the United States is asleep at the wheel—before, all at once, the government shifts into gear in the most extraordinary fashion. There will be a moment—in just a few years, just a couple more “2023-level” leaps in model capabilities and AI discourse—where it will be clear: we are on the cusp of AGI, and superintelligence shortly thereafter. While there’s a lot of flux in the exact mechanics, one way or another, the USG will be at the helm; the leading labs will (“voluntarily”) merge; Congress will appropriate trillions for chips and power; a coalition of democracies formed.

Startups are great for many things—but a startup on its own is simply not equipped for being in charge of the United States’ most important national defense project. We will need government involvement to have even a hope of defending against the all-out espionage threat we will face; the private AI efforts might as well be directly delivering superintelligence to the CCP. We will need the government to ensure even a semblance of a sane chain of command; you can’t have random CEOs (or random nonprofit boards) with the nuclear button. We will need the government to manage the severe safety challenges of superintelligence, to manage the fog of war of the intelligence explosion. We will need the government to deploy superintelligence to defend against whatever extreme threats unfold, to make it through the extraordinarily volatile and destabilized international situation that will follow. We will need the government to mobilize a democratic coalition to win the race with

authoritarian powers, and forge (and enforce) a nonproliferation regime for the rest of the world. I wish it weren't this way—but we will need the government. (Yes, regardless of the Administration.)

In any case, my main claim is not normative, but descriptive. In a few years, The Project will be on.

The path to The Project

A turn-of-events seared into my memory is late February to mid-March of 2020. In those last weeks of February and early days of March, I was in utter despair: it seemed clear that we were on the covid-exponential: a plague was about to sweep the country, the collapse of our hospitals was imminent—and yet almost nobody took it seriously. The Mayor of New York was still dismissing Covid-fears as racism and encouraging people to go to Broadway shows. All I could do was buy masks and short the market.

And yet within just a few weeks, the entire country shut down and Congress had appropriated trillions of dollars (literally >10% of GDP). Seeing where the exponential might go ahead of time was too hard, but when the threat got close enough, existential enough, extraordinary forces were unleashed. The response was late, crude, blunt—but it came, and it was dramatic.

The next few years in AI will feel similar. We're in the midgame now. 2023 was already a wild shift. AGI went from a fringe topic you'd be hesitant to associate with, to the subject of major Senate hearings and summits of world leaders. Given how early we are still, the level of USG engagement has been impressive to me. A couple more "2023"s, and the Overton window will be blown completely open.

As we race through the OOMs, the leaps will continue. By 2025/2026 or so I expect the next truly shocking step-changes; AI will drive \$100B+ annual revenues for big tech companies and outcompete PhDs in raw problem-solving smarts. Much as

the Covid stock-market collapse made many take Covid seriously, we'll have \$10T companies and the AI mania will be everywhere. If that's not enough, by 2027/28, we'll have models trained on the \$100B+ cluster; full-fledged AI agents/drop-in remote workers will start to widely automate software engineering and other cognitive jobs. Each year, the acceleration will feel dizzying.

While many don't yet see the possibility of AGI, eventually a consensus will form. Some, like Szilard, saw the possibility of an atomic bomb much earlier than others. Their alarm was not well-received initially; the possibility of a bomb was dismissed as remote (or at least, it was felt that the conservative and proper thing was to play down the possibility). Szilard's fervent secrecy appeals were mocked and ignored. But many scientists, initially skeptical, started realizing a bomb was possible as more and more empirical results came in. Once a majority of scientists came to believe we were on the cusp of a bomb, the government, in turn, saw the national security exigency as too great—and the Manhattan Project got underway.

As the OOMs go from theoretical extrapolation to (extraordinary) empirical reality, gradually, a consensus will form, too, among the leading scientists and executives and government officials: we are on the cusp, on the cusp of AGI, on the cusp of an intelligence explosion, on the cusp of superintelligence. And somewhere along here, we'll get the first genuinely terrifying demonstrations of AI: perhaps the oft-discussed "helping novices make bioweapons," or autonomously hacking critical systems, or something else entirely. It will become clear: this technology will be an utterly decisive military technology. Even if we're lucky enough to not be in a major war, it seems likely that the CCP will have taken notice and launched a formidable AGI effort. Perhaps the eventual (inevitable) discovery of the CCP's infiltration of America's leading AI labs will cause a big stir.

Somewhere around 26/27 or so, the mood in Washington will become somber. People will start to viscerally feel what is happening; they will be *scared*. From the halls of the Pentagon to the backroom Congressional briefings will ring the obvious

question, the question on everybody's minds: do we need an AGI Manhattan Project? Slowly at first, then all at once, it will become clear: this is happening, things are going to get *wild*, this is the most important challenge for the national security of the United States since the invention of the atomic bomb. In one form or another, the national security state will get very heavily involved. The Project will be the necessary, indeed the only plausible, response.

Of course, this is an extremely abbreviated account—a lot depends on when and how consensus forms, key warning shots, and so on. DC is infamously dysfunctional. As with Covid, and even the Manhattan Project, the government will be incredibly late and hamfisted. After Einstein's letter to the President in 1939 (drafted by Szilard), an Advisory Committee on Uranium was formed. But officials were incompetent, and not much happened initially. For example, Fermi only got \$6k (about \$135k in today's dollars) to support his research, and even that was not given easily and only received after months of waiting. Szilard believed that the project was *delayed for at least a year* by the short-sightedness and sluggishness of the authorities. In March 1941, the British government finally concluded a bomb was inevitable. The US committee initially entirely ignored this British report for months—until finally in December 1941, a full-scale atomic bomb effort was launched.

There are many ways this could be operationalized in practice. To be clear, this doesn't need to look like literal nationalization, with AI lab researchers now employed by the military or whatever (though it might!).¹¹³ Rather, I expect a more suave orchestration. The relationship with the DoD might look like the relationship the DoD has with Boeing or Lockheed Martin. Perhaps via defense contracting or similar, a joint venture between the major cloud compute providers, AI labs, and the government is established, making it functionally a project of the national security state. Much like the AI labs "voluntarily" made commitments to the White House in 2023, Western labs might more-or-less "voluntarily" agree to merge in the national effort. And likely Congress will have to be involved, given the trillions of investment involved, and for checks-and-balances.¹¹⁴ How all these details shake out is a story for another day.

¹¹³ Note that while private companies help develop components for nuclear weapons, they are never allowed to possess a completed and assembled nuclear weapon. In comparison, the mainline version of the "AGI government project" I am putting forward here is unprecedentedly privatized, for the WMD reference class.

¹¹⁴ Congress—even the Vice President!—didn't know about the Manhattan Project. We probably shouldn't repeat that here; I'd even suggest that key officials for The Project require Senate confirmation.

But by late 26/27/28 it will be underway. The core AGI research team (a few hundred researchers) will move to a secure location; the trillion-dollar cluster will be built in record-speed; The Project will be on.

Why The Project is the only way

I AM UNDER no illusions about the government. Governments face all sorts of limitations and poor incentives. I am a big believer in the American private sector, and would almost never advocate for heavy government involvement in technology or industry.

I used to apply this same framework to AGI—until I joined an AI lab. AI labs are very good at some things: they've been able to take AI from an academic science project to the commercial big stage, in a way only a startup can. But ultimately, AI labs are still startups. We simply shouldn't expect startups to be equipped to handle superintelligence.

There are no good options here—but I don't see another way. When a technology becomes this important for national security, we will need the USG.

Superintelligence will be the United States' most important national defense project

I've discussed the power of superintelligence in previous pieces. Within years, superintelligence would **completely shake up the military balance of power**. By the early 2030s, the entirety of the US arsenal (like it or not, the bedrock of global peace and security) will probably be obsolete. It will not just be a matter of modernization, but a wholesale replacement.

Simply put, it will become clear that the development of AGI will fall in a category more like nukes than the internet. Yes, of course it'll be dual-use—but nuclear technology was dual-use too. The civilian applications will have their time. But in

the fog of the AGI endgame, for better or for worse, national security will be the primary backdrop.

We will need to completely reshape US forces, within a matter of years, in the face of rapid technological change—or risk being completely outmatched by adversaries who do. Perhaps most of all, the initial priority will be to deploy superintelligence for defensive applications, to develop countermeasures to survive untold new threats: adversaries with superhuman hacking capabilities, new classes of stealthy drone swarms that could execute a preemptive strike on our nuclear deterrent, the proliferation of advances in synthetic biology that can be weaponized, turbulent international (and national) power struggles, and rogue superintelligence projects.

Whether nominally private or not, the AGI project will need to be, *will be*, integrally a defense project, and it will require extremely close cooperation with the national security state.

A sane chain of command for superintelligence

The power—and the challenges—of superintelligence will fall into a very different reference class than anything else we’re used to seeing from tech companies. It seems pretty clear: this should not be under the unilateral command of a random CEO. Indeed, in the private-labs-developing-superintelligence world, it’s quite plausible individual CEOs would have the power *to literally coup the US government*.¹¹⁵ Imagine if Elon Musk had final command of the nuclear arsenal.¹¹⁶ (Or if a random nonprofit board could decide to seize control of the nuclear arsenal.)

It is perhaps obvious, but: as a society, we’ve decided democratic governments should control the military;¹¹⁷ superintelligence will be, at least at first, the most powerful military weapon. The radical proposal is not The Project; the radical proposal is taking a bet on private AI CEOs wielding military power and becoming benevolent dictators.

(Indeed, in the private AI lab world, it would likely be even worse than random CEOs with the nuclear button—part of AI

¹¹⁵ It wouldn’t even require cooperation from AI lab employees at this point, since they’ll have been mostly automated by this point.

¹¹⁶ And as Sam Altman once said, every year we get closer to AGI everybody will gain +10 crazy points.

¹¹⁷ In fact, the government having the biggest guns was an enormous **civilizational achievement!** Rather than medieval-like fights of all against all, we would sort out disagreements via courts, pluralistic institutions, and so on.

labs' abysmal security is their utter lack of internal controls. That is, random AI lab employees (with zero vetting) could go rogue unnoticed.)

We will need a sane chain of command—along with all the other processes and safeguards that necessarily come with responsibly wielding what will be comparable to a WMD—and it'll require the government to do so. In some sense, this is simply a Burkean argument: the institutions, constitutions, laws, courts, checks and balances, norms and common dedication to the liberal democratic order (e.g., generals refusing to follow illegal orders), and so on that check the power of the government have withstood the test of hundreds of years. Special AI lab governance structures, meanwhile, collapsed the first time they were tested. The US military could already kill basically every civilian in the United States, or seize power, if it wanted to—and the way we keep government power over nuclear weapons in check is not through lots of private companies with their own nuclear arsenals. There's only one chain of command and set of institutions that has proven itself up to this task.

Again, perhaps you are a true libertarian and disagree normatively (let Elon Musk and Sam Altman command their own nuclear arsenals!).¹¹⁸ But once it becomes clear that superintelligence is a principal matter of national security, I'm sure this is how the men and women in DC will look at it.

The civilian uses of superintelligence

Of course, that doesn't mean the civilian applications of superintelligence will be reserved for the government.

- The nuclear chain reaction was first harnessed as a government project—and nuclear weapons permanently reserved for the government—but civilian nuclear energy flourished as private projects (in the 60s and 70s, before environmentalists shut it down).
- Boeing made the B-29 (the most expensive defense R&D project during WWII, more expensive than the Manhat-

¹¹⁸ Or perhaps you say, just open-source everything. The issue with simply open sourcing everything is that it's not a happy world of a thousand flowers blooming in the US, but a world in which the CCP has free access to US-developed superintelligence, and can outbuild (and apply less caution/regulation) and take over the world. And the other issue, of course, is the proliferation of super-WMDs to every rogue state and terrorist group in the world. I don't think it'll end well. It's a bit like how having no government at all is more likely to lead to tyranny (or destruction) than freedom.

In any case, people overrate the importance of open-source as we get closer to AGI. Given cluster costs escalating to hundreds of billions, and key algorithmic secrets now being proprietary rather than published as they were a couple years ago, it'll be 2-3 or so leading players, rather than some happy community of decentralized coders building AGI.

I do think a different variant of open source will continue to play an important role: models that lag a couple years behind being open sourced, helping the benefits of the technology diffuse broadly.

tan Project) and the B-47 and B-52 long-range bombers in partnership with the military—before using that technology for the Boeing 707, the commercial plane that ushered in the jet era. And today, while Boeing can only sell stealth fighter jets to the government, it can freely develop and sell civilian jets privately.

- And so it went for radar, satellites, rockets, gene technology, WWII factories, and so on.

The initial development of superintelligence will be dominated by the national security exigency to survive and stabilize an incredibly volatile period. And the military uses of superintelligence will remain reserved for the government, and safety norms will be enforced. But once the initial peril has passed, the natural path is for the companies involved in the national consortium (and others) to privately pursue civilian applications.

Even in worlds with The Project, a private, pluralistic, market-based, flourishing ecosystem of civilian applications of superintelligence will have its day.

Security

I've gone on about this at length in a previous piece in the series. On the current course, we may as well give up on having any American AGI effort; China can promptly steal all the algorithmic breakthroughs and the model weights (literally a copy of superintelligence) directly. It's not even clear we'll get to "North Korea-proof" security for superintelligence on the current course. In the private-startups-developing-AGI-world, superintelligence would proliferate to dozens of rogue states. It's simply untenable.

If we're going to be at all serious about this, we obviously need to lock this stuff down. Most private companies have failed to take this seriously. But in any case, if we are to eventually face the full force of Chinese espionage (e.g., stealing the weights being the MSS's #1 priority), it's probably impossible for a private company to get good enough security. It will require

extensive cooperation from the US intelligence community at that point to sufficiently secure AGI. This will involve invasive restrictions on AI labs and on the core team of AGI researchers, from extreme vetting to constant monitoring to working from a SCIF to reduced freedom to leave; and it will require infrastructure only the government can provide, ultimately including the physical security of the AGI datacenters themselves.

In some sense, security alone is sufficient to necessitate the government project—both the free world’s preeminence and AI safety are doomed if we can’t lock this stuff down. (In fact, I think it’s fairly likely to be a major factor in the ultimate trigger: once the Chinese infiltration of the AGI labs becomes clear, every Senator and Congressperson and national security official will... have a strong opinion on the matter.)

Safety

Simply put: there are a lot of ways for us to mess this up—from ensuring we can reliably control and trust the billions of superintelligent agents that will soon be in charge of our economy and military (the superalignment problem), to controlling the risks of misuse of [new means of mass destruction](#).

Some AI labs claim to be committed to safety: acknowledging that what they are building, if gone awry, could cause catastrophe and promising that they will do what is necessary when the time comes. I do not know if we can trust their promise enough to stake the lives of every American on it. More importantly, so far, they have not demonstrated the competence, trustworthiness, or seriousness necessary for what they themselves acknowledge they are building.

At core, they are startups, with all the usual commercial incentives. Competition could push all of them to simply race through the intelligence explosion, and there will at least be some actors that will be willing to throw safety by the wayside. In particular, we may want to “spend some of our lead” to have time to solve safety challenges, but Western labs will need to coordinate to do so. (And of course, private labs will have already had their AGI weights stolen, so their safety pre-

cautions won't even matter; we'll be at the mercy of the CCP's and North Korea's safety precautions.)

One answer is regulation. That may be appropriate in worlds in which AI develops more slowly, but I fear that regulation simply won't be up to the nature of the challenge of the intelligence explosion. What's necessary will be less like spending a few years doing careful evaluations and pushing some safety standards through a bureaucracy. *It'll be more like fighting a war.*

We'll face an insane year in which the situation is shifting extremely rapidly every week, in which hard calls based on ambiguous data will be life-or-death, in which the solutions—even the problems themselves—won't be close to fully clear ahead of time but come down to competence in a “fog of war,” which will involve insane tradeoffs like “some of our alignment measurements are looking ambiguous, we don't really understand what's going on anymore, it might be fine but there's some warning signs that the next generation of superintelligence might go awry, should we delay the next training run by 3 months to get more confidence on safety—but oh no, the latest intelligence reports indicate China stole our weights and is racing ahead on their own intelligence explosion, what should we do?”.

I'm not confident that a government project would be competent in dealing with this—but the “superintelligence developed by startups” alternative seems much closer to “praying for the best” than commonly recognized. We'll need a chain of command that can bring to the table the seriousness that making these difficult tradeoffs will require.

Stabilizing the international situation

The intelligence explosion and its immediate aftermath will bring forth one of the most volatile and tense situations mankind has ever faced. Our generation is not used to this. But in this initial period, the task at hand will not be to build cool products. It will be to somehow, desperately, make it through this period.

We'll need the government project to win the race against the

authoritarian powers—and to give us the clear lead and breathing room necessary to navigate the perils of this situation. We might as well give up if we can't prevent the instant theft of superintelligence model weights. We will want to bundle Western efforts: bring together our best scientists, use every GPU we can find, and ensure the trillions of dollars of cluster build-outs happen in the United States. We will need to protect the datacenters against adversary sabotage, or outright attack.

Perhaps, most of all, it will take American leadership to develop—and if necessary, enforce—a nonproliferation regime. We'll need to subvert Russia, North Korea, Iran, and terrorist groups from using their own superintelligence to develop technology and weaponry that would let them hold the world hostage. We'll need to use superintelligence to harden the security of our critical infrastructure, military, and government to defend against extreme new hacking capabilities. We'll need to use superintelligence to stabilize the offense/defense balance of advances in biology or similar. We'll need to develop tools to safely control superintelligence, and to shut down rogue superintelligences that come out of others' uncared-for projects. AI systems and robots will be moving at 10-100x+ human speed; everything will start happening extremely quickly. We need to be ready to handle whatever other six-sigma upheavals—and concomitant threats—come out of compressing a century's worth of technological progress into a few years.

At least in this initial period, we will be faced with the most extraordinary national security exigency. Perhaps, nobody is up for this task. But of the options we have, The Project is the only sane one.

The Project is inevitable; whether it's good is not

ULTIMATELY, MY MAIN CLAIM here is descriptive: whether we like it or not, superintelligence won't look like an SF startup, and in some way will be primarily in the domain of national security. I've brought up The Project a lot to my San Francisco friends in the past year. Perhaps what's surprised me most is

how surprised most people are about the idea. They simply haven't considered the possibility. But once they consider it, most agree that it seems obvious. If we are at all right about what we think we are building, *of course*, by the end this will be (in some form) a government project. If a lab developed literal superintelligence tomorrow, *of course* the Feds would step in.

One important free variable is not *if* but *when*. Does the government not realize what's happening until we're in the middle of an intelligence explosion—or will it realize a couple years beforehand? If the government project is inevitable, earlier seems better. We'll dearly need those couple years to do the security crash program, to get the key officials up to speed and prepared, to build a functioning merged lab, and so on. It'll be far more chaotic if the government only steps in at the very end (and the secrets and weights will have already been stolen).

Another important free variable is the international coalition we can rally: both a tighter alliance of democracies for developing superintelligence, and a broader benefit-sharing offer made to the rest of the world.

- The former might look like the **Quebec Agreement**: a secret pact between Churchill and Roosevelt to pool their resources to develop nuclear weapons, while not using them against each other or against others without mutual consent. We'll want to bring in the UK (Deepmind), East Asian allies like Japan and South Korea (chip supply chain), and NATO/other core democratic allies (broader industrial base). A united effort will have more resources, talent, and control the whole supply chain; enable close coordination on safety, national security, and military challenges; and provide helpful checks and balances on wielding the power of superintelligence.
- The latter might look like **Atoms for Peace**, the **IAEA**, and the **NPT**. We should offer to share the peaceful benefits of superintelligence with a broader group of countries (including non-democracies), and commit to not offensively using superintelligence against them. In exchange, they refrain from pursuing their own superintelligence projects, make

safety commitments on the deployment of AI systems, and accept restrictions on dual-use applications. The hope is that this offer reduces the incentives for arms races and proliferation, and brings a broad coalition under a US-led umbrella for the post-superintelligence world order.

Perhaps the most important free variable is simply whether the inevitable government project will be competent. How will it be organized? How can we get this done? How will the checks and balances work, and what does a sane chain of command look like? Scarcely any attention has gone into figuring this out.¹¹⁹ Almost all other AI lab and AI governance politicking is a sideshow. This is the ballgame.

The endgame

AND SO BY 27/28, the endgame will be on. By 28/29 the intelligence explosion will be underway; by 2030, we will have summoned superintelligence, in all its power and might.

Whoever they put in charge of The Project is going to have a hell of a task: to build AGI, and to build it fast; to put the American economy on wartime footing to make hundreds of millions of GPUs; to lock it all down, weed out the spies, and fend off all-out attacks by the CCP; to somehow manage a hundred million AGIs furiously automating AI research, making a decade's leaps in a year, and soon producing AI systems *vastly* smarter than the smartest humans; to somehow keep things together enough that this doesn't go off the rails and produce rogue superintelligence that tries to seize control from its human overseers; to use those superintelligences to develop whatever new technologies will be necessary to stabilize the situation and stay ahead of adversaries, rapidly remaking US forces to integrate those; all while navigating what will likely be the tensest international situation ever seen. They better be good, I'll say that.

For those of us who get the call to come along for the ride, it'll be . . . stressful. But it will be our duty to serve the free

¹¹⁹ To my Progress Studies brethren: you should think about this, this will be the culmination of your intellectual project! You spend all this time studying American government research institutions, their decline over the last half-century, and what it would take to make them effective again. Tell me: how will we make The Project effective?



Figure 38: Oppenheimer and Groves.

world—and all of humanity. If we make it through and get to look back on those years, it will be the most important thing we ever did. And while whatever secure facility they find probably won't have the pleasantries of today's ridiculously-overcomped-AI-researcher-lifestyle, it won't be so bad. SF already feels like a peculiar AI-researcher-college-town; probably this won't be so different. It'll be the same weirdly-small circle sweating the scaling curves during the day and hanging out over the weekend, kibitzing over AGI and the lab-politics-of-the-day.

Except, well—the stakes will be all too real.

See you in the desert, friends.

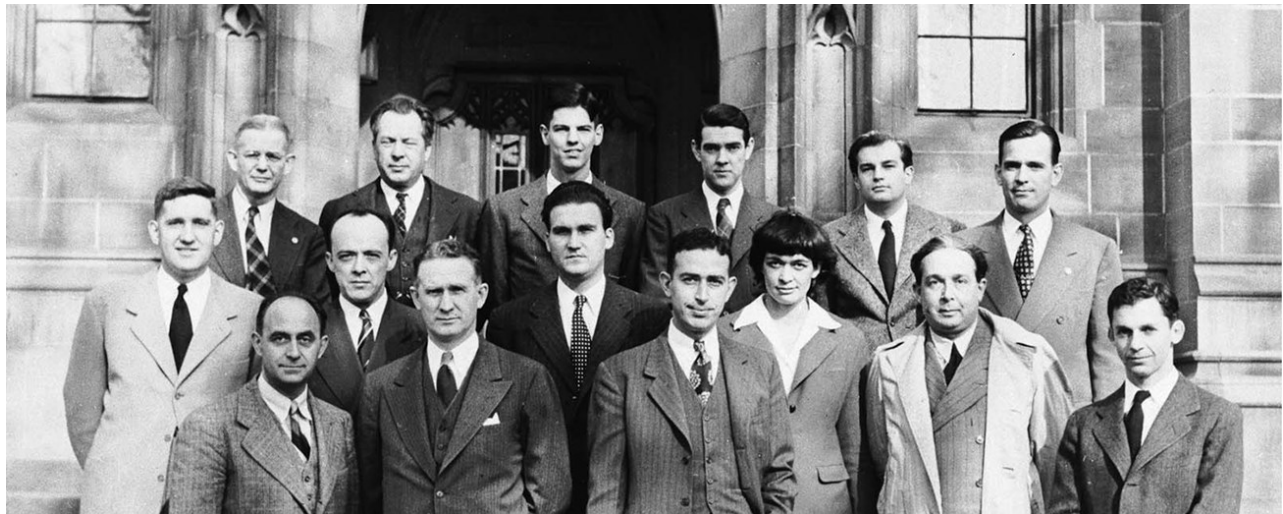


Figure 39: Reunion of atomic scientists on the fourth anniversary of the first controlled nuclear fission reaction at UChicago.