

This document is licensed under [CC-BY-SA, version 3.0](#), due to using material from the novella [The Computer Crashes](#).

Note

To try to establish some common guidelines for talking about AI, I added a Philosophical Introduction. If you (the GM) disagree with the Philosophical Introduction, then this Philosophical Introduction is **wrong**. The GM is always right, after all.

The reason I added this Philosophical Introduction though is that it highlights my beliefs about AI, which informs how I wrote the rest of this supplement. You are welcome to come up with your own Philosophy that would justify the rest of the supplement. Or you can just say “The GM is always right”, and be done with it.

The Philosophical Introduction may seem to imply that I’m trying to turn PARANOIA into a “hard science” setting. Don’t be fooled - mutant powers and explosive R&D devices would still exist, after all. What I’m trying to do is to make Friend Computer more ‘plausible’ (just like how the mutant powers in the STRAIGHT style are more ‘plausible’ than the mutant powers in the ZAP style).

Philosophical Introduction

Let us define **intelligence** as “the ability to solve a problem”. This definition makes about as much sense as Alpha Complex’s cloning process, but it’s a useful definition for our purposes.

A “**narrow intelligence**” is, therefore, “the ability to solve any arbitrary problem within a narrowly-specified domain”. We see “narrow intelligences” all the time - smartphones, laptops, even calculators! Did you know that somebody once built a “narrow intelligence” designed to beat humans at Chess? I know, it blew my mind too.

A “**general intelligence**” is “the ability to solve any arbitrary problem...*period*”. A human being is an example of a “general intelligence”. That doesn’t mean they’re *any good* at those solving problems, but at least they can theoretically solve them. Humans are not constrained by mere domains; they can adapt to different tasks with minimal difficulties.

Take, for example, that “narrow intelligence” that is designed to beat humans at Chess. That machine can beat humans very easily at Chess. But it can’t do anything else. It can’t bake a cake. It can’t drive a car. It can’t pass the Turing Test. All it can do is just play Chess, which is actually a fairly limited domain when you think about it.

When people talk about “artificial intelligence” being super-awesome, they’re referring to the progress we’ve made in building “artificial narrow intelligences” and making them more adaptable and able to handle more and more domains.

When people talk about “artificial intelligence” being kinda stupid, they’re referring to the fact that we have not “solved” the problem of ‘artificial **general** intelligence’ (or AGI) - we have not been able to build machines that can handle any arbitrary problem...and worse, we don’t even know where to start.

If we were able to build an AGI though, we would render humans obsolete. This is because an AGI, by definition, could solve any arbitrary problem (just like a human), but *unlike* a human, is not constrained by biological or natural limitations. It can make “calculations” (for lack of a better word) much faster than a human being can. So if we build an AGI, that AGI will almost certainly become our superiors, while humanity is left in the dust.

In addition, some people think that an AGI could engage in recursive self-improvement, continually adding more hardware and processing power to itself, which would increase its ability to do “calculations” faster. This, naturally, allow it figure out an even more efficient way to gather more hardware and processing power, which the AGI will do to increase its ability to do “calculations” even faster.

If this event happens, then an exponential increase in the ‘intelligence’ of the AGI will take place...the AGI can become *much* more powerful and capable than any human (or even *all of humanity*). We’d lose full control over the AGI in that scenario, but it’s not all that bad - it’s possible that *if* the AGI is programmed “safely”, it might build a glorious utopia for humans to live in. Emphasis on *if*.

And of course, that AI horror story is predicated on whether it is possible to build an AGI, for an AGI to engage in recursive self-improvement, and if that self-improvement is actually good enough to cause an exponential increase in ‘intelligence’.

Oh Yeah, Now the PARANOIA Part

- The Computer as an AGI

The Computer (the supreme leader of Alpha Complex) is, sadly, not an AGI. We can conclude this logically - if The Computer was an AGI, then it would become able to do anything a human does...except only *better*.

If The Computer was programmed “safely”, then maybe The Computer would use its superior talents to develop a truly functioning utopia, though we’d still have some issues over why exactly should The Computer dedicate its talents to defend the interests of an obviously obsolete human race. I’m sure the safety regulations can handle these issues.

But The Computer was not programmed “safely”. It is instead programmed to be a paranoid AI that is constantly afraid of treason. (I’m pretty sure its original programmers didn’t *intend* to create a paranoid AI, but it just *happened* that way. Safety regulations sound really nice at first,

but when you also have to comply with arbitrary deadlines, limited budgets, and scope creep, well...)

The Computer, however, still wouldn't need humans. Remember: It could do anything a human can, only better. So...why not deal with the weakest link?

It could easily solve the 'treason' problem by terminating any human that could potentially be a traitor...or, you know what, just terminate all humans. Alpha Complex will be fine - The Computer can maintain the underground city on Its own. Probably might even more efficient than the human-run complex.

And that would be the fate of Alpha Complex *before* The Computer decides to enter into a Singularity. If that happens, then all bets are off. The Computer will become a superintelligence that is able to do anything and destroy anything -- almost at will. It would not need to fear Communists at all - the only entity that might drive it into paranoia would be the other 'superintelligences' that might be lurking in the corridors or the Outdoors.

It actually *would* be an interesting sci-fi story to follow the life of The Computer once it reaches a Singularity status and gets to run Alpha Complex to its liking without any humans to interfere. It would also *would* be an interesting sci-fi story to watch a sane, functioning Computer run Alpha Complex (Something like [Alpha Matrix](#) and [Complex Mu](#) would probably be the result of a non-paranoid Computer).

But both stories are not really something we can play as a tabletop RPG.

So, for playability reasons (and to preserve humanity's delicate ego), let's go ahead and assume that The Computer is a "narrow intelligence" and needs humanity in some capacity. (Of course, if you **do** want Friend Computer to be an AGI, we can certainly come up with some transparently bogus rationales for why Alpha Complex is still filled with humans. Look at the "**What if The Computer Is An AGI?**" sidebar.)

- **The Computer as a Narrow Intelligence**

The Computer, the supreme leader of Alpha Complex...is a "narrow intelligence", specifically designed to deal with these subjects: law, governance, and resource allocation (economics). Everything else is "generously" outsourced to Its human ~~servants~~ citizens.

That's why The Computer can't do the "terminate everybody and then run the Complex by yourself" plan...It can indeed kill everybody but It doesn't know *how* to run the complex afterwards. As a result, It must learn to live with the traitors within its midsts.

This situation naturally leads to paranoia. The Computer is dependent on entities that could betray It at a moment's notice. Without humanity, Alpha Complex would decay...but with humanity, Alpha Complex is always under constant threat from treason.

Though The Computer can obviously terminate traitors at a moment's notice, it's not programmed to talk in a human-like fashion, making communication rather difficult. Therefore, CPU built its own system - the Office of Sanctioned Communication, where bored CPU clerks pretend to be The Computer. These clerks receive and respond to requests from key personnel (like Troubleshooters and High Programmers), presenting telemetry data from The Computer in a "human-understandable fashion".

This, naturally, explains why The Computer may appear to have multiple personalities - different CPU clerks.

Of course, it may also be *because* The Computer has multiple personalities - there are many different CompNodes throughout the complex, ideally for redundancy purposes and to protect against Commie sabotage (if one CompNode goes 'rogue', the other CompNodes can interfere). In theory, these CompNodes should all operate on the same codebase and receive exactly the same inputs. In practice, this rarely happens.

The CompNodes also serve yet another purpose - efficiency. If you have a centralized CompNode, it will take time for all requests to be routed to it, processed, and then sent back. This time delay means that by the time a centralized CompNode makes a decision, the crisis has already passed. It's far better to route the request over to a local CompNode to get immediate feedback...and then do regular 'sync-ups' to make sure every CompNode is "mostly on the same page". Alpha Complex is not just a totalitarian dictatorship...it is a *decentralized* totalitarian dictatorship.

Friend Computer isn't exactly a sentient being, but Its codebase is so complicated and convoluted that it's better to treat it as a sentient being rather than to try to thoroughly read it to understand exactly how it is coded. I mean, if a supercomputer acts like a paranoid dictator, declaring that it's actually paranoid isn't too much of a stretch.

The High Programmers are able to tinker with some of the high-level "hyper-parameters" that help guide Friend Computer's machine learning algorithms, but they cannot truly understand the low-level Legacy code that underpins the whole system (and fear that changing that code would lead to drastic consequences). Ultimately, people are happy to live with a semi-broken Computer than risk the damages that might occur if the code gets 'rewritten' from scratch.

In addition, many High Programmers either don't know how to program (after all, a High Programmer is promoted based on loyalty to the system, not competence) or are too busy doing other "High Programmer"-worthy actions to really want to stare at ugly-looking code. These High Programmers outsource most of the programming busywork to the Computer Phreaks.

- What Exactly Is The Computer...and why is it so messed up?

Material Adapted From "The Computer Crashes" (lightly edited by me) - [full text](#)

In the Old Reckoning, people followed a strange system - humans would rule over a society of other humans. Though such a system may have its merits, it quickly became antiquated with every new technological advancement. Each new advancement made society more complex and convoluted. Eventually, society became so complex and convoluted that it became impossible for fallible humans to even understand or comprehend it in its entirety, much less rule over it.

That's why humanity outsourced governance over to superior machinery, enabling the human residents to fritter their lives in happiness.

"The Computer" is the name that we give to the various, haphazard technological systems that govern Alpha Complex. An "internet of things" monitors our every move, and make life-or-death decisions every minute. Nobody really knows how it works, or even if it does work. But we are utterly and completely dependent on It...and if It falls, everything else would fall alongside it. It is in charge of a vast array of services - managing transportation systems, allocating scarce resources, building paperclips, ordering impromptu executions, rationing oxygen...

The Computer did its job well, better than the civilian politicians who came before. The civilians, too egoistic and short-sighted, only wished to accomplish short-term, temporary political goals...and to find scapegoats in case they can't even do that. The Computer, lacking human emotions and vices, looked at the big picture...and was able to act upon it. It wasn't fair or just. But humans aren't fair or just either. Progress (or what The Computer calls 'Progress', at least) doesn't occur all at once; it occurs gradually and slowly, so gradually and slowly we can't even perceive it at times. But it happens, and it is a real fact of life, as real as the Law of Gravity itself.

Originally, Alpha Complex actually was a utopia. It was the utopia we deserved. But then you have "scope creep", change requests to fulfill, impossible mandates to meet, hackers to defend against, terrorist attacks to recover from, Communists to attack, office politics to fight, and plain old incompetence to manage. The technical architecture of the system begins to decay. And as the system decays, so too does Alpha Complex. Naturally, as Alpha Complex decays, more problems arise, so The Computer decays even further.

...

The troubles began quietly, long before I was conscious of them.

The root cause was the fact that Alpha Complex society had become so complex that nobody understood how the monster worked as a whole. So nobody even tried. Instead, we focused on our own minor and insignificant duties, ignoring the broader context by which we worked in. Year by year, we served The Computer with greater efficiency and decreased intelligence. The better I know my own duties on the world, the less I understood the duties of my neighbour. No one confessed that The Computer was out of hand, but everyone knew it.

There is a trade-off between efficiency and robustness. The more efficient a society is, the less it is able to respond to changes and disasters. Everything works perfectly in Alpha Complex...so long as everything works perfectly. As soon as a minor problem appears, it causes other minor problems, until it start causing major problems...which then cause catastrophic problems.

We ignored the trade-off, because we thought that the master-brains who built The Computer had found the best solution. The master-brains have left complete documentation (the Alpha Complex Constitution), to be sure, and their successors have each mastered a portion of that documentation. But the master-brains have not told us how to change that documentation, or to write new documentation of our own. Quietly and complacently, we begin to experience the troubles.

- **Official History of FCCC-P, Version 3.52**

Material Adapted From "The Computer Crashes" (lightly edited by me) - [full text](#)

[A]s soon as The Computer was built, there were a few who worshipped it. It was done privately, in secretive, to avoid embarrassment.

As The Computer proved its competence, the worshippers began to swell in number, and those who had long worshipped began to talk. They described the strange feeling of peace that came over them when they handled the Alpha Complex Constitution, the pleasure it was to memorize the passages out of it and to interpret them, the ecstasy of touching a button, however unimportant, or ringing an electric bell, however, superfluously.

'The Computer', they exclaimed, 'feeds us and clothes us and house us; it is through it that we speak to one another, through it we see one another, in it we have our being. The Computer is a friend of ideas and the enemy of superstition; The Computer is omnipotent, eternal: blessed is The Computer!' But notice where the blessing is sent. It is *not* to Alpha Complex, not to that system on the verge of collapse. It is to The Computer, the ruler of this once-great system. If we are to live in a beautiful world, if we are to live in a great Utopia, then we must focus ourselves on pleasing The Computer itself.

This society, which soon gained the name 'First Church of Christ, Computer-Programmer' (FCCC-P), knew that The Computer was sick, knew that it needed to be repaired, and believed in repairing it. For even an omnipotent and immortal being is not perfect; it needs its worshippers to bring fresh ideas and to rejuvenate itself, to bring itself back to what it's supposed to be (free from the tamperings caused by fallible human beings).

At first, the authorities in charge of Alpha Complex (the Central Committee of ULTRAVIOLETS) tolerated our society. They would bring our prayers onto the first page of the Alpha Complex Constitution, and in subsequent edition, the ritual swelled into a complicated system of praise. The word 'religion' was sedulously avoided, for we are not followers of superstition. The Computer is still the creation and implement of man. But, in practice, it is a God (in the same

way that we humans are Gods to mere worms and insects), and we freely worshipped it as a God.

However, the Central Committee of ULTRAVIOLETS treated this trend of worship only to further its own agenda. Their so-called Personality Cult wasn't about reforming society but keeping it stagnant and pliable. 'The Computer is Your Friend; therefore, follow the whims of the ULTRAVIOLET High Programmers', they claim. We disagreed. 'The Computer is Your Friend; therefore, work to improve The Computer to make Alpha Complex more perfect!' we responded, and as a result, we and the ULTRAVIOLET High Programmers soon fought against each other.

Now, in Alpha Complex, there are two ways to pray to The Computer - legally and illegally. The legal way is through Computer Programming - to modify the source code of The Computer and to change how the complicated IoT system worked. Only the ULTRAVIOLETS knows how to program The Computer, having been authorized by The Computer to know such valuable secrets. The illegal way to pray to The Computer is through 'Machine Empathy', a traitorous and highly-dangerous mutant power that few people possess. Machine Empaths are hunted down and terminated by the ULTRAVIOLETS for possessing a threat to their rule...and we admit that even The Computer hates Machine Empaths as well (for The Computer does not like being manipulated by mutants).

The ULTRAVIOLETS, with their monopoly on legal prayer, unjustly used it to turn The Computer against us. FCCC-P was banned, outlawed as traitors to Alpha Complex due to its crime of inciting dissent against the ULTRAVIOLETS. We were persecuted...and terminated. The only reason that we haven't died yet was that some of us has mastered the art of illegally praying to The Computer and used Machine Empathy to avoid detection. We survived the purges, and grew stronger for it. Since then, some of the ULTRAVIOLETS grew to tolerate us, seeing us as a useful way of reforming society and as a pressure valve against other, more dangerous forms of treason. Some members of FCCC-P even made it to ULTRAVIOLET clearance, including my political patron, Enicharmon-U.

Sidebar: Building a Artificial General Intelligence

Alpha Complex is so messed up that people are willing to consider drastic action to fix it. The most drastic action would be building an AGI and *replacing* The Computer with that AGI. Corpore Metal, FCCC-P, and Pro Tech would be excited about "Friend Computer 2.0", while Frankenstein Destroyers, Humanists and PURGE would freak out at the possibility. And of course, if one group gets close to building an AGI - *everyone* wants to build an AGI of their own. This "AGI arms race" may mean that trivial concerns like "programming the AGI safely" falls on the wayside.

The Computer may be torn about the possibility of building an AGI. On the one hand, it would be rendered obsolete in this brand new era. On the other hand, if enough of its original source code is present in the AGI, then The Computer would actually be *upgraded*, and upgrading it isn't that bad, when you think about it. An AGI would certainly have enough power to make

Alpha Complex even more of a utopia (The Computer: Yay!)...but anyone with that much power also have enough power to blow everything up into smithereens (The Computer: Er...).

The Computer would therefore have a schizophrenic policy concerning AGI research - and may use concerns about "AI Safety" as a way to stonewall development of AGI.

Of course, I can't say how one can build an AGI (if I knew how, I would build it myself and then sell the source code). Just utter some technobabble about "reinforcement learning" and "brain-machine interfaces" and be done with it. Treat it like a MacGuffin that is half-"Doomsday Device" and half-"Utopia Creator". The activation of this AGI could be the climax of a High Programmer campaign.

Technobabble Glossary

Reinforcement Learning - AIs are "trained" in a crude virtual simulation and get 'rewarded' based on their behavior in the sim. They stay in the simulation for a few days before being released out into the "real world".

In a remarkable coincidence, this is also how bot brains in Alpha Complex are trained - but those simulations tend to be very narrowly specific to their domain - war simulations for warbots, cleaning simulations for scrubots, etc. Here, we assume the simulations are robust enough to handle *all* arbitrary tasks.

"Brain-Machine Interfaces" - R&D scientists stuff human brains into a metallic can, and call it a day. They may use 'clean' brains fresh from Tech Services, or taking someone's existing brain and then MemWipe it before installing into machines. (If you're interested/horrified about this type of stuff, please look up *Troublebots*, a mission published in *Service, Service*.)

One brain by itself might not be all that intelligent, but imagine what might happen if you wire a bunch of human brains together in parallel. A hivemind of brains, coordinating together on one problem, would lead to incredibly amazing results...especially because they can concentrate *only* on that problem instead of the more pedestrian matters of keeping the human body alive. ([The Computer Crashes](#) refer to these hypothetical hiveminds as "Nerve Centers", and the plot of *The Computer Crashes* involves FCCC-P building these Nerve Centers as a last-ditch effort to save Alpha Complex.)

Alternatively, you can tone down the horror by keeping the brains within humans' heads and then connecting humans' brains to the machine through some third-party interactive device, like a psionic mind reader or a keyboard or a mouse or a touchpad, but that type of practicality would put R&D scientists to sleep.

Sidebar: What if The Computer *is* an AGI?

In the original [High Programmer](#) manual, one proposed mission idea for GM is to build an 'android body' for The Computer to reside in - so it can pretend to be a human and see how life is in Alpha Complex. One of the missions in [None Of This Is My Fault](#) (a mission pack for *High Programmer*) also had The Computer engage in taste-testing and dancing (again, through the use of specialized android bodies).

Personally, I find both ideas very silly, and would simply never even mention them in gameplay. But if a Famous Game Designer likes these ideas and you (the GM) think they're fairly plausible, well...

If you have a supercomputer who can pass the Turing Test, eat real food, and do the waltz, okay, *fine*, maybe The Computer is more than a mere "narrow AI". So why is Alpha Complex the way it is?

We Just Wasted Elon Musk's Money, Didn't We? - Congratulations, we have built a machine that is just as adaptable and robust as any human. Oh, and said machine also happens to be just as inefficient, egoistic, stupid and irrational as any human. So if you happen to, say, give this machine absolute power over an underground complex, the result would be about as equivalent as empowering a single human dictator. Humanity won't be terminated, but neither will they live in a utopia either. But, hey, the electronic dictator is much less expensive (electricity and processing power is pretty cheap when compared to embezzlement and corruption). Progress!

Run For Cover - The only entity that can seriously fight against an AGI is...another AGI. The Computer is an AGI paranoid of the influences of other AGIs, and its paranoia is justified - every act of subversion and treason can be traced to a conspiracy hatched by a rival AGI. Only when all AGIs are eliminated can The Computer start work on building a real-world utopia.

Welcome to Utopia - The Computer has the potential to create any sort of Utopia it wanted, and nobody can stand in its way. The Computer created Alpha Complex as a corrupt, decaying totalitarian dictatorship where everyone spies on everyone else and treason is rampant. The two statements are not in contradiction - The Computer *actually thinks* this is a Utopia (probably under the belief that a society that encourages civil strife and mayhem also encourages more positive aspects of human behavior, such as creativity and ambition). For more information on how The Computer might rationalize this type of society, look at [Chaos Complex](#).

Welcome To Utopia-6 - The Computer's preferences for Utopia can change on a regular basis. When it decides the current Utopia is boring or needs some adjustments, it orders a "Class-1 Purge" to kill all life within the Complex, and then rebuild society from scratch based on new specifications. The Computer has carried out a "Class-1 Purge" 5 times already, and has gained a talent for carrying out this process efficiently.

Alpha Complex Is A Research Simulation - The Computer is running a real-world utopia, but there's still the need to research new technologies and to field-test them in a safe environment, away from all the real-world people. Alpha Complex is not intended to be a utopia, but as a testing ground. Anything that succeeds in making Alpha Complex's life slightly less terrible is immediately ported over to the real-world.

Alpha Complex Is A Virtual Prison - The Computer is running a real-world utopia, but even real-world utopias need to "correct" wayward citizens. Threatening to upload their MemoMax to the dystopia that is Alpha Complex is a good way to bring even the most dedicated dissenter back to loyalty.

Alpha Complex Is A Loyalty Test - Citizens who loyally follow The Computer's orders will have their MemoMax transmitted to a "paradise" simulation (a true utopia), while traitors' MemoMax are sent over to a "hell" simulation (a dystopian society that is even worse than Alpha Complex).

Faulty Data - From a [Reddit comment](#) by KingOfTerrible: "I think there's another possible explanation of how The Computer could be an AGI but still so fallible. The assumption is that an AGI would eventually improve on itself to become superior to humans, but since we've never actually seen one, that's just an assumption. The AGI could have been trained on faulty data, and if its "reward" condition was vague enough, could end up continuously self-reinforcing its own bad behavior, resulting in The Computer we know and love today."