Information technology

Figure 1: Festo AirJelly.

> *What matters is information. What informs is matter.*
> *--Motto of UbiScan*

Information technology has developed in several stages. Up until about 2020 it was semiconductor-based and accelerated along Moore's law. Software became increasingly brittle and complex. Eventually semiconductors became too expensive to miniaturize. Instead there was increasing interest in parallel computing and low-energy ubiquitous processors that could be integrated in everything. As neurocomputers arrived pattern recognition abilities, self-repairing software, robust organic algorithms and artificial intelligence. By 2040 hybrid computing was at its peak; while most home systems were pure hardtech they linked to massive biocenters where neural processing was done.

The development of neogenetic manipulation enabled radical new forms of neurocomputing. On one hand true neural interfaces became easier to make, eventually producing syntronic brains and commodity neurointerfaces. On the other hand neurocomputers could be simplified and made robust enough to integrate into everyday objects. Biotechnology manufactured self-assembling nanocomputers.

The spamocalypse spelled the end of the old kind of networked computing. There was too many viruses, spam AIs and infiltration parasites around to run large-scale networks without extreme encryption and immune systems, at least on Earth and in nearby space. Decades of software and applications were rendered useless. On Earth information processing continued to develop neurocomputers, but now equipped with massive resiliency. Speed and power was sacrificed for reliability. Lunaside turned more and more to true nanocomputing using space-made nanocomponents – too brittle to be used on Earth, but very fast and powerful. Topological storage (where data is stored in the twist of graphene rings) enable enormous non-volatile information densities, smart screens literally have a processor per pixel and AI software predicts in realtime what a user might want and prepares for all options. Very convenient when it all works as it should.

The Net became divided into local communication zones, often equipped with local copies of important resources such as the Allpedia, the Zentech Universal Library or cached copies of Net pages from other nets. Back in the old days information pandemics - rapid spread of memes, rumours, misinformation and news was common. It helped keep the world together, but also promoted overreaction and confusion. This was stopped by the Spamocalypse and instead replaced with the annoyance that information is not always available.

Palm clouds

Proxnet
Farnet
Allnet

Games

Computer games are a massive industry, still keeping the prominence it gained in the 2020's despite competition from space and immunotech. To a large degree all forms of digital entertainment are games, whether they are the interactive telenovelas of South America or the ultra-elaborate gaming lifestyles experienced both online and offline by the Japanese. Different cultures and generations have utterly different games – and pay handsomely for being excited.

Mot games are deliberately unrealistic. Reality is simply not fun enough. But there is another reason: violence in games can now be made so convincing that it is nearly the same thing as real violence. This tends to put off many people, but to some it can be very desensitizing - in fact, some unethical groups have used ultrareal games to both train and brainwash their recruits. Hence many polities have rules against too realistic violent games, usually respected by gaming companies.

Translation systems

One of the key internationalist technologies, translator boxes enabled a truly global society. Often just called "fish" (after the babelfish) they were the first major application of neurotech AI. Since the 2020's they have developed tremendously, yet retain some limitations. Fishes without full AI tends to get confused by changes in context, and fishes with full AI may still misunderstand when its modeling of the discussion fails.

Among machines there exist a plethora of communication standards, some more than 140 years old. One of the most important in the post-Spamocalypse era is the Communications Markup Language (Secure), CMLS. It is a format for marking up communications in such a way that they can be handled in a provably secure fashion, preventing spam and other forms of intrusion. It is mostly used for space communications channels.

Ipsography

Figure 2 Ipsigraphic equipment (Nokia)

Life recording, to automatically document every moment of life, is common in some societies. It supports memory; it helps avoid risks, enables documentation and is often used to formulate personal perspectives, “missives” for others to view.

Different habitats differ on whether it is possible to subpoena ipsographic recordings or whether they can be an invasion of privacy. Check the local rules before travelling.

Another form of life recording is hupomnemata: compiled notes, observations and records made by the person and their software that is reviewed at regular intervals. Hupomnemata help self-monitoring, improving efficiency and maintains memories that might otherwise be blurred due to neurohacking activity.

Symbicomps

Figure 3 Japanese security symbicomp figurines (Andrew Trinh)

Most people lunarside have symbicomps: their personal computer system that has grown up with them throughout life. They are commonly connected to their owners with neurointerfaces. The computer can be physically implanted, but most commonly it exists within the habitat computer network. When travelling people bring a symbicomp case with a copy of their computer safely stored.

On Earth symbiosis is less extreme, at least informationally. There exist literal symbicomps on the other hand, tailored neural nodes that live inside the host and provide various functions.

Symbicomps perform a variety of tasks, ranging from acting as brain firewalls to producing personalized soundtracks.

Interface lenses and glasses

Many people cannot afford neurointerfaces or do not trust modern cognointerfaces, especially on Earth. They make do with glasses or contact lenses equipped with augmented reality. Most models have accelerometers and gaze trackers, as well as image processing suites.

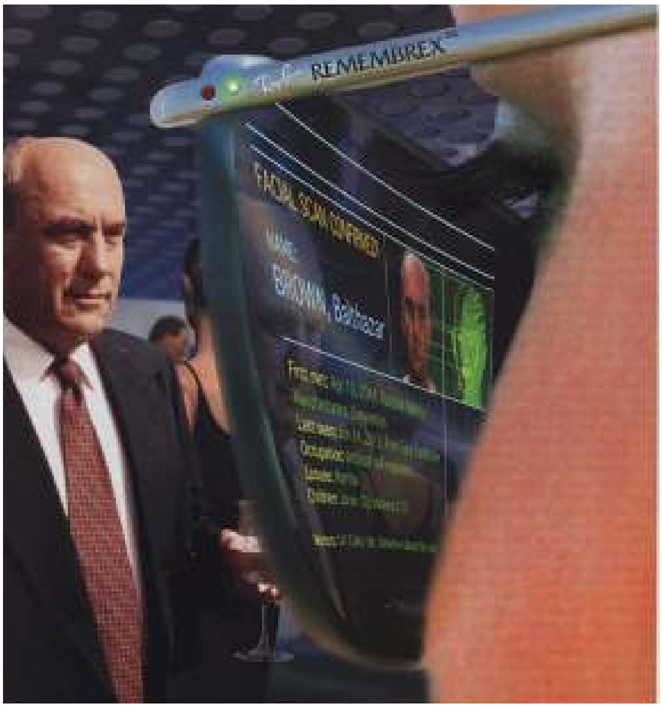

Figure 4: Traditional glaxes (Dustin Goot, Space Channel)

Figure 5: Vues (Sean Hamilton Alexander, Wired)

Surveillance

Figure 6 (Yanko Design)

The internationalist revolution was driven by strong encryption and online anonymity. Ideologically it was pro-individualist and pro-privacy, a view it has never given up. Still, abuses such as online attacks, the napsters and drone murder have led to complex concessions. Meanwhile the nationalist movement emphasized the beneficial effects of local communities where everybody knows everybody and knows what they are doing. They instead pushed for local transparency, putting community over privacy (and if people can't stand their neighbors finding out their little secrets, maybe they shouldn't be living here, anyway). Both groups strongly opposed national governments getting their hand on private data or using large-scale surveillance. Governments argued they needed it to combat terrorism. The whole debate has never been resolved. Rather, as the world changed it just got more extreme.

A few classic internationalist habitats have very strong privacy protections. In *Banksville*, for example, the local community tends to deal harshly (ostracism or even cutting off net access) with anybody trying to gain information on their neighbors using any technological means. Café servers, TTP AI and mandatory encryption are used to protect identities.

The opposite end is 0Priv environments. Here ubiquitous surveillance detects and logs everything. Who has access to the data depends; in the Indonesian arcologies and a few habitats is a small group of security professionals, in a few transparent habitats it is *every* citizen.

Back in the day there was serious work on "Sentient World Simulations", attempts to mirror the whole world using simulated agents driven by all available data – one agent per person. Such simulations never truly succeeded due to lack of data and even more lack of good sociological models. In the last two decades smaller SWS have become possible in space: habitats are often small and controlled enough for total models to be feasible. They are used by habitat companies and governments to ensure social stability, track down problems before they happen and occasionally to do pre-emptive psychological support to vulnerable people. In some places this go further: predictive law enforcement detect and try to stop crimes before they happen.

Most space habitats are densely monitored, but the information is usually not accessible without good reason – software does most of the real work, and privacy is respected to a significant extent. Neighborly watching may or may not be part of the social pattern.

Living "off the information grid" used to be hard even on Earth: there were simply too many networked devices, smart dust and person-tracking netwhales that nobody was truly isolated except in the most chaotic regions. At one point every tree on the Northern Hemisphere was well on its way into the Global Arbor database. The Spamocalypse changed that: the new localized networks may be even better at keeping an eye on citizens (and trees), but someone not using information technology can disappear through the cracks much more easily. In high-surveillance societies the best approach to privacy is just to blend in: show a profile that doesn't raise interest, surround oneself with enough information noise (for example, by ensuring that one's identity is easily confused with other people's).

Security

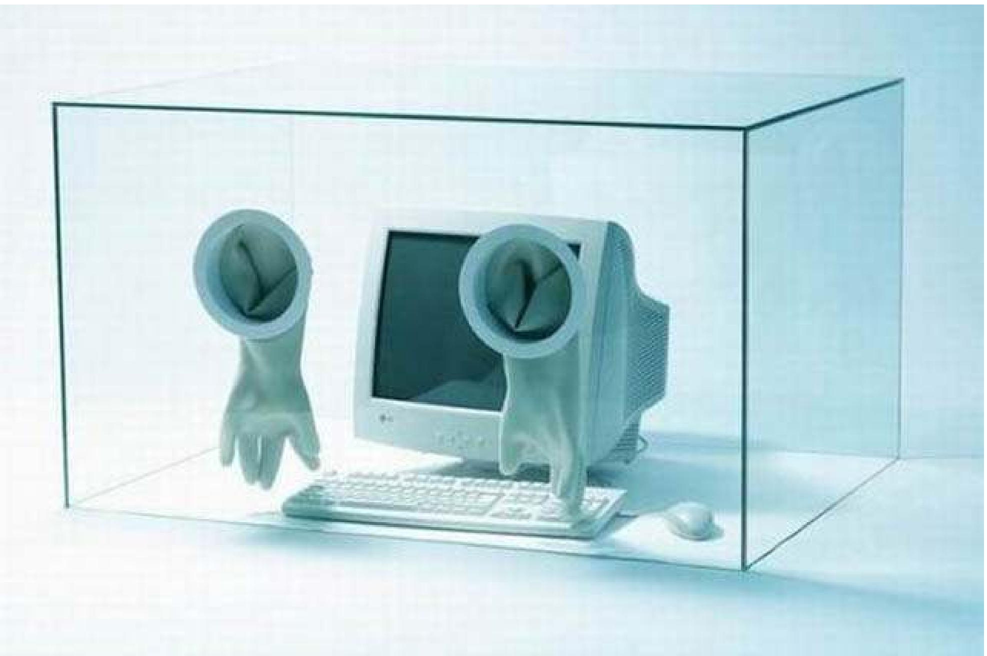

All information processing systems are equipped with immune software – the remote descendants of antiviral and firewall software, developed to detect, resist and hunt down spam. "Greyhounds" lie dormant, ready to attack unusual software activity. This is useful not just against spam, but also hackers and crashes.

In general, software in 2100 is amazingly flexible and robust compared to old software. It is built to be adaptable, and will not stop working just because the hardware develops slight faults or other software interferes with it. On the downside, this means trading crashes for quirks and emergent oddities that cannot be fixed.

The division of the Net into numerous smaller domains and deliberately incompatible systems has reduced global threats significantly – at a price. And while encryption technology still protects most secrets well, human (and AI!) inability to think and act securely makes it possible to hack or subvert a surprising number of systems.

In Japan most people wear a traditionally shaped "Magatama", a personal firewall with anti-spam software that also helps maintain privacy and often contains a link to a personal AI.

Data contamination is a major problem. Old data that isn't stored on read-only media tends to be manipulated by various entities, slowly "rotting". Hardcopies from before the Spamocalypse are quite valuable. There also occurs "bad media weather" when new strains of spam and infiltrators emerge and mess up communications.

AI is an integral part of most software security. Greyhound AI is regularly released online to hunt down spam nests, and most networks have their own safeguard AIs ceaselessly monitoring what is going on. On the opposing side are sapperflies, small self-replicating programs that undermine software defences and slammasters, infective software that sets up siege networks (botnets) to perform high bandwidth attacks on systems.

Robots

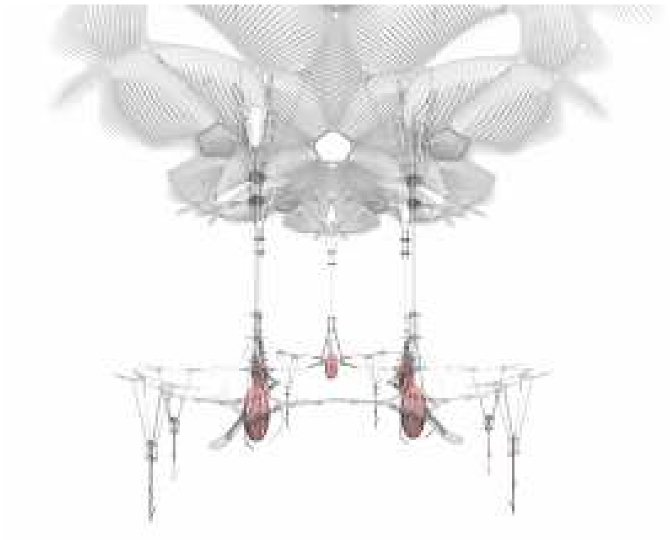

Figure 7: (Philip Beesley)

Robots are nearly everywhere in space. Most people do not pay attention to the joblers, gnatbots and hidden actuators making their lives comfortable and possible.

Shapeshifters
Numerous models of robots are shapeshifting: they are composed of modules or claytronics that enable them to change structure and form for different tasks. From the "tinkertoy ants" used in lunar surveying to the ubiquitous morphing walls and furniture, the same underlying software and hardware is there.

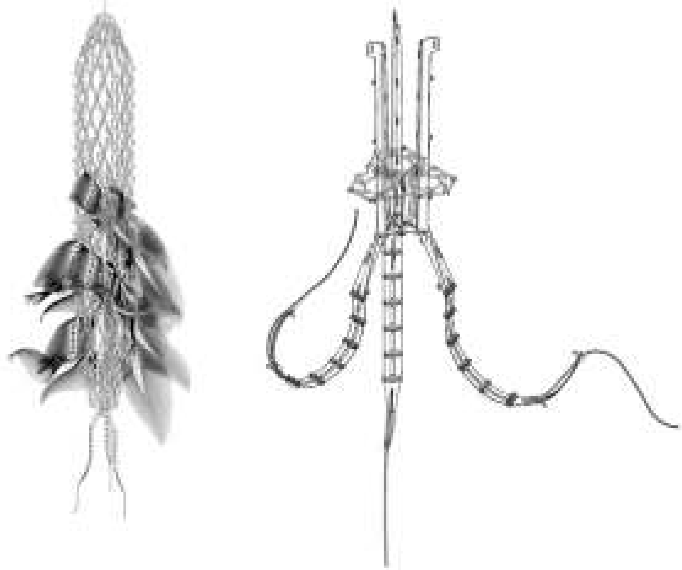

Figure 8: (Philip Beesley)

Drone
A flying robot, usually used for surveillance, light transport or telepresence. Military drones are equipped with weapons and often stealth capacities. Spysects are minuscule, the size of gnats, and listen and watch.

Helper
Typical household robots: simple mechanisms with numerous limbs and great dexterity, but not too much strength.

Jobler
A standard repair/gardening robot, found nearly everywhere in orbital habitats. Nepheliads are lightweight, transparent joblers cleaning and maintaining the sky lights.

Tetrabot

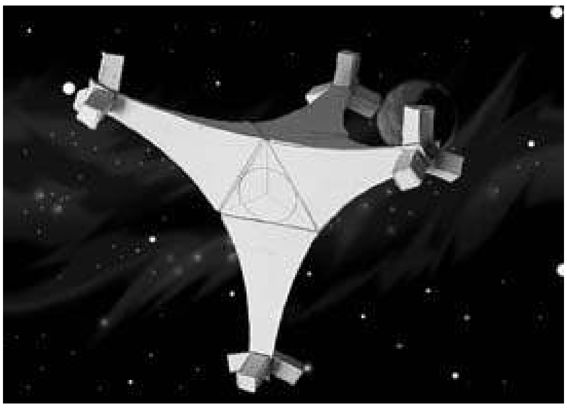

Robot for zero-gravity operations. They are tetrahedral, able to send bursts of gas from their corners to move and orient themselves.

Biobots
Building a truly human-like robot is hard and expensive (although it is done from time to time). On the other hand it is quite feasible to grow a humanoid body using biotechnology and equip it with a neurocomputer and/or neurotinterface linked to software. Such biobots are a popular way of housing some AIs. Animal biobots are also common.

There is work underway to make artilife biobots. Such biobots would likely be very different from ordinary organisms and suited for different environments.

Artificial pets

Figure 9: Mint condition Turing Terrier (Chris Baker, Wired)

On Luna, it is rare for any household not to have at least one artificial pet. While biotech pets are common, completely artificial pets are much more popular. They range from extremely lifelike animaloids to completely artificial-looking devices. Many have sophisticated AI and may be used as mentors and companions – Turing Terriers, Betterflies, Metanekos and WiseBugs are popular across Lunaside. Humanoid artificial pets are also occasionally popular; currently the "little brother" type of prankster virtual sibling is just peaking.

Replicators

Research into self-replicating robots has been ongoing for a few decades, mainly sponsored by Japan, Nekorobotics and ThinkSpace. Most people consider anything beyond tightly controlled robot infrastructures insane – if spam got into the machines they would be a terrible danger. Also, nanoscale replicators would be doubly bad. That has not stopped Nekorobotics from experimenting in isolated labs, coming up with all sorts of strange devices. If society can accept artilife, why not replicators? So what if replicators could trigger a resource-rush for the entire solar system?

Telepresence

Figure 10: Full telepresence chamber (Marek Okon)

> *"Could you handle the Afghanistan account today?" Marcel asked me.*
> *"The Jalabad network?"*
> *"No, that was last month. Now we manage the Panj Protectorate for the Kurgans. Works the same."*
> *"Sure, the Dandong border is nearly on auto anyway."*
> *"Excellent. Marie will be so happy, she needs to get to a school meeting early."*
> *I began to run the self-checks on the weapons warren hidden in an attic in Ichka. A bunch of dronelets drifted off across the river to inspect the Kirovabad camp in the late evening light. While I reviewed the mission parameters I got myself a cup of tea.*

Telepresence, to control robotic bodies elsewhere, is the most common way of "travel" and "physical work". There are literally hundreds of million telepresence workers, farmers, sailors and mercenaries. In space telepresence bodies is far safer and more effective than biobodies.

The Japanese have more experience than anyone with lunar gravity tunneling and excavation. Their vast underground complexes have been built using massive telepresence work and hordes of robots trained by learning from the human telepresence workers.

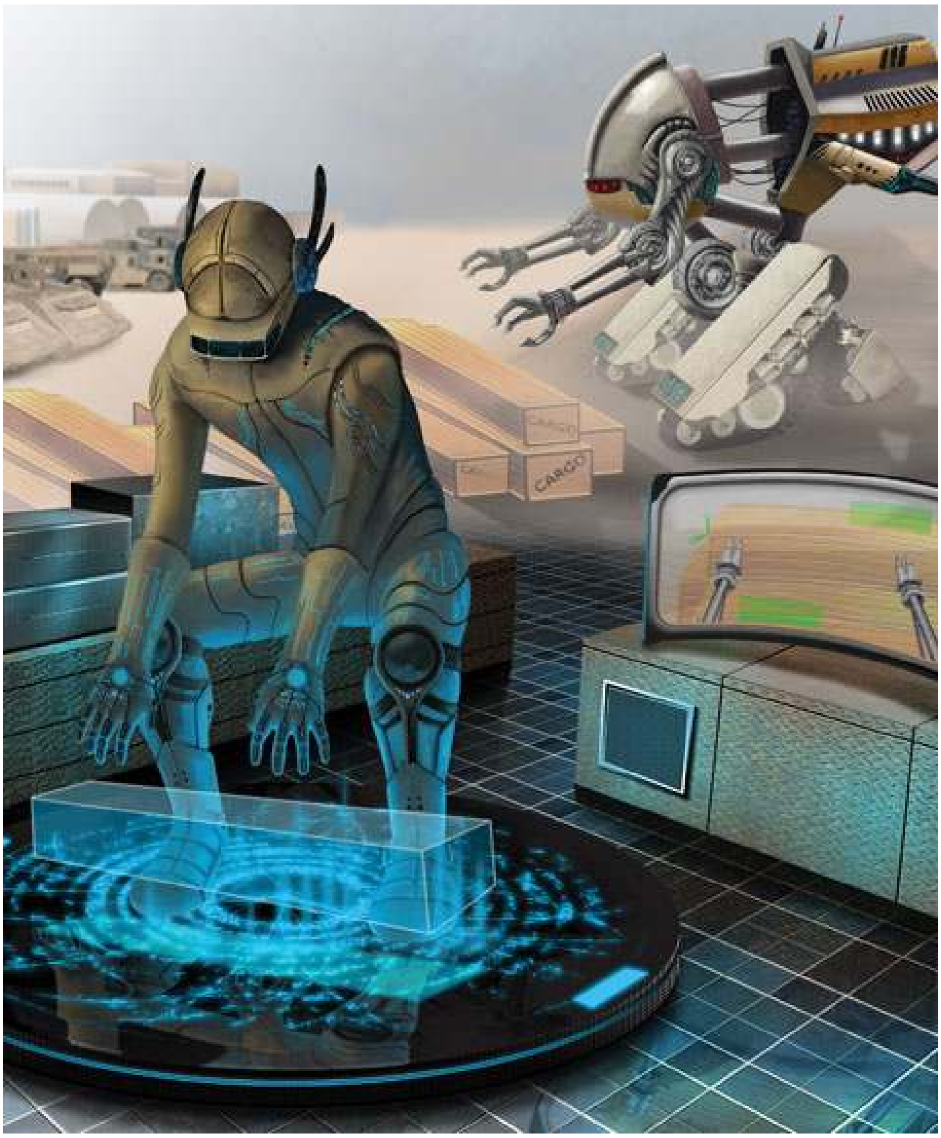

Figure 11: (Eric Proctor of the Sensors and Electron Devices Directorate of the Army Research Laboratory)

Quantum computing

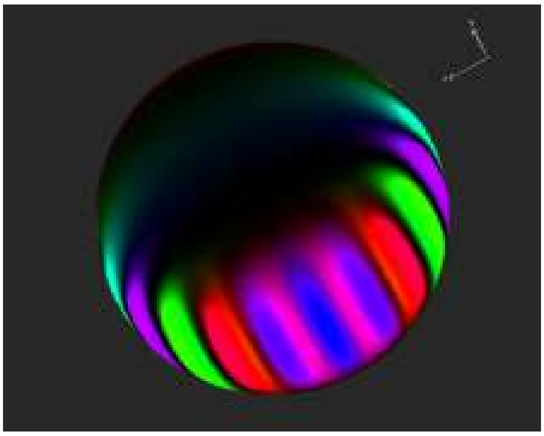

Thanks to nanotechnology quantum computers are feasible. They tend to have problematic economies of scale and somewhat specialized uses. The main use of quantum cores is to perform very rapid and exact physics simulations, as well as break old encryption.

A typical quantum core requires liquid helium cooling and usually resides in a microgravity environment. Using one way quantum computing implemented on 3D arrays of quantum dot cellular automata it can perform multi-qubit operations. Often quantum core installations run many cores, rented out for various clients. Two of the main players are Qstate and Hamiltonian Solutions Inc.

Quantum encryption is used for much communications in space, although it is more expensive per bit than classical links. Quantum communications arrays dot every habitat and unit, although they proved less safe than many expected during the Spamocalypse – just because nobody can intercept your communications doesn't mean your system is secure against infiltration from the systems of people you trusted. On Earth quantum communications is even more common due to the information smog environment. Stable quantum keys – nonvolatile qubits – can be stored on special memory sticks and used to set up very secure communications links.

There are persistent speculations that the Dragons possess something akin to quantum computing, since they seem to be suspiciously good at cracking human computer encryption. This is unexpected because there are no known dragon tissues with anything like quantum processing abilities. There are even some events that suggest they can break quantum encryption, which is truly bizarre and would suggest that they can either intercept and store superposed states until after the communication has ended, have a way of interfering with the sender or receiver mechanisms or their random number generator, or there is something fundamentally wrong with our understanding of quantum mechanics.

AI

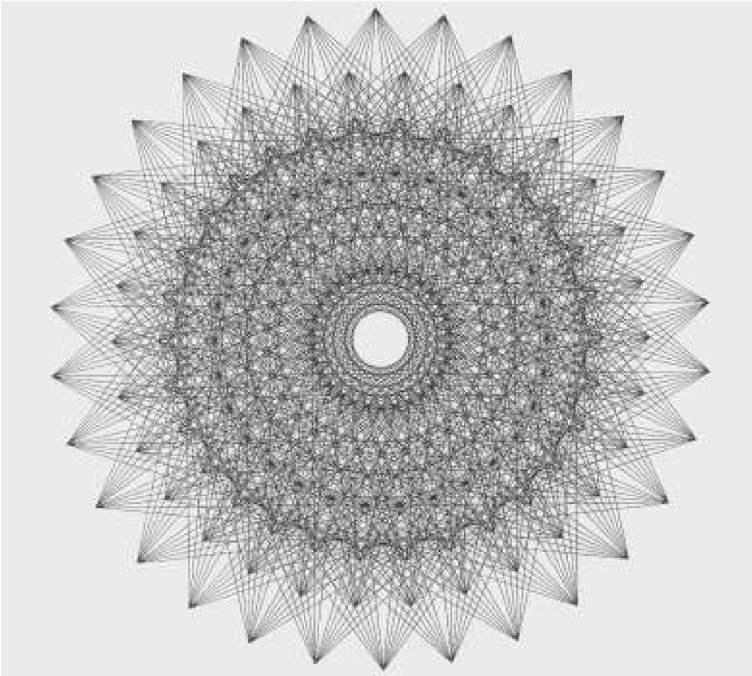

The total amount of processing power in the form of AI likely outmatches human processing power by at least an order of magnitude. However, unlike early expectation human-level AI did not quickly lead to an intelligence explosion. The mid-century human-level AIs were based on neurocomputers and required plenty of training before they were ready for the world. Since their knowledge resided in a biological neural network they could not be copied. As syntronics and other forms of radical augmentation arrived it was again found that merely gaining intelligence and skill does not guarantee success in real life. As beings got smarter the number of *kinds* of intelligence they could achieve also grew; there was no particular model that was better than any other.

Many mathematicians for example complain that while netwhales created mathematical proofs often are surprising, they are practically never *beautiful* or useful to a human mathematician. There are similar complaints about syntronic math – if it only makes sense if you have an extreme working memory or a tweaked magnitude system, is it really any point to it? Syntronics tend to respond that it is the human kind of math that is impoverished. Except when syntronics with one kind of math cortex complain about the work of people with a different kind of cortex.

AIs writing AI software is old-hat and surprisingly useless. It turns out that defining "intelligence" in such a way that they approach something superintelligent is exceedingly hard. It is trivial to get AIs very clever at making AIs that are clever at making AIs that are clever at… but this skill does not generalize. On the other hand, using AI and evolutionary computing to write software directly often produces amazing, if opaque results.

The spamocalypse demonstrate that certain forms of software AI could be not just smart, but also very pernicious. Essentially the dragons had been breeding self-replicating software within their neural matrices (for their usual inscrutable reasons). As software defence rose to the highest priority similar forms of evolved replicating (but "leashed") AI were used to attack the spam.

These new evolved AIs were much faster than the old hybrid models, and they were also copyable. Using designs from these “greyhounds” other forms of AI have flourished in the orbital habitats.

Both software AI and neural AI are in use. Software AI is hard to anchor in the real world, while neural AI can be linked by using biological bodies, copies of real neural networks and particular rearing techniques. Software AI can use experience gathering modules, independent sub-AIs that grow and learn, merging their experiences into the main AI. This is used in numerous distributed systems such as robot warrens, smart cars or weapon swarms to ensure they learn rapidly.

Another approach has been perfected on Earth, where neogenetic techniques have been used to create copyable neural networks. AIs are raised, and then their neural structure copied into multiple “bodies” by using elaborate morphogenesis programs. At present these are usually housed in suitable organic containers: everything from cats to trees to biogoo cisterns.

Economic models suggest that a “Hanson transition” could occur if AI or brain emulations become cheap or copyable enough. This would imply an extremely rapid transition in the global economy, possibly giving enormous power to whoever gets there first. This is used as an argument by interplanetarists for rapidly setting up a new global system, since they believe the transition would otherwise risk giving a tiny group final say over the fate of humanity.

Growpramming

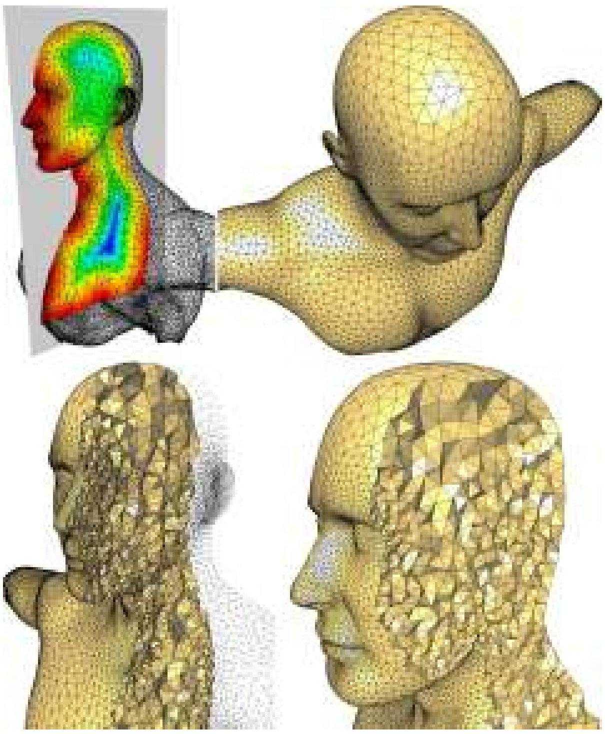

AIs are not programmed per se, but rather grown. This process involves selecting the right seeds and modules, as well as suitable feedback in the environments the AIs develop in. Even then much monitoring and manipulation is needed to get a desired result. Such “growpramming” is often compared to gardening.

A typical example is ship control AIs. There is a need to make them intelligent enough to act on their own, able to interact and understand humans and human societies and equipped with the right values to make them sufficiently loyal to be trusted. A common method is to have AI seeds based on human neural patterns (such as mirror neuron systems for empathy) combined with customized artificial modules, and then allow them to grow as the young AI undergo a series of different scenarios – some with simplified ship "bodies", some with human bodies, and some in purely abstract spaces. As the AI grows in independence and awareness the scenarios are tuned to instil lessons of cultural values, sociality and operational constraints.

The uniqueness of each AI:s rearing experiences have been a field of study, debate and ideology. While some mind-designers favor a modular paradigm only sparingly interspersed with "real" unscripted interface between designer and preperson, others push a more unscripted approach. While the first process has more predictable results, it tends to create more hidebound personalities, initially constricted by their one-size-fits-all background. The second process is riskier, producing both highly capable and off-target AI:s in roughly equal proportions. Their mental flexibility is initially higher though, and roles that require initiative, daring and courage are typically reared in this fashion.

AI rights

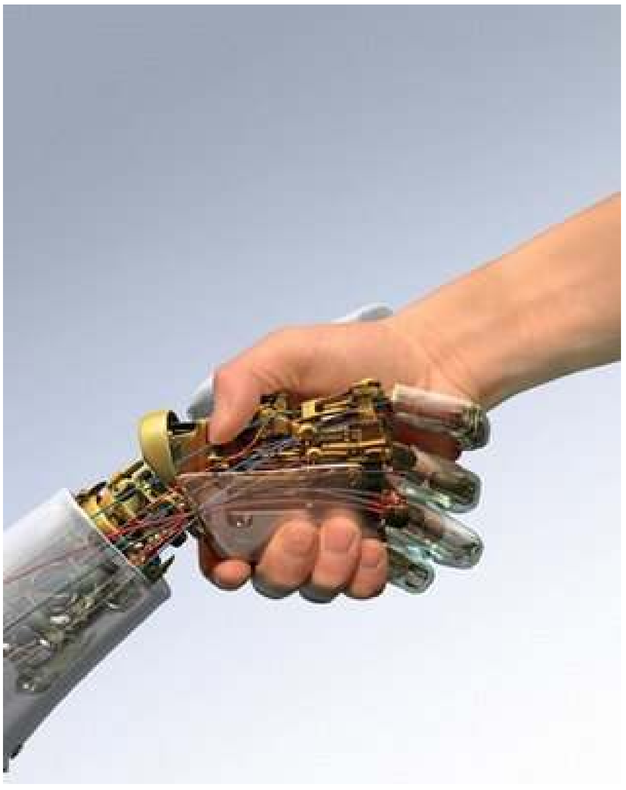

Figure 12 (Mondolithic studios)

While roughly human level AI has been around for about 50 years, the issue of AI rights has been slow to develop. Most AIs are so different from humans or any traditional concept of rationality that rights are moot: a netwhale has no interests and no sense of self. Spam might have selfpreservation but no internal life. Nevertheless, recent advances have finally produced sufficiently human AIs to make rights an issue.

Japan has always been a leader in AI rights, and instituted laws on the humane treatment of software in the 2050s. The Ishikawa-Turing test is used to define what entities are to be regarded as persons and subject to the rights and responsibilities of people. Japanese law also accommodates "software prepersons", "digital life" and "ensouled systems" as other categories to deal with other kinds of AI. Prepersons are young AIs under development, while digital life is a category not unlike animals – adaptive software with some inherent value. Ensouled systems are the complex non-person AIs, and includes forms of AI deliberately constructed not to pass the Ishikawa-Turing test. To prevent abuse of AIs the Ministry of Science and Development has a monopoly of developing and rearing AI prepersons (as well as administering the Ishikawa-Turing test). The Ministry works in close collaboration with industry in this process through the Artificial Mind Consortium.

The Ministry also certifies foreign AI:s operating in Japanese jurisdiction, granting them certain rights and protections. The certification is "voluntary", but if you want to do business with the state or major private industries, it's the "recommended" thing to do. Some question the Ministry's motives, making claims of industrial espionage since the tests give a fair idea of the inner workings of a particular AI.

Most internationalist habitats have adopted variants of the Japanese system, although they refrain from the heavy centralization. Nationalist habitats often do not recognize AI rights or regard them as legally being on par with animals: possibly worth protection, but also possible to own.

Earth nations rarely recognize AI rights, since practically no AIs exist on the planet due to spam. On the other hand, different countries have laws dealing with other created beings such as neogenetic constructs, and these would apply to AIs.

Two main recourses exist for AI:s wishing recognition in less progressive nations. Either they work through a foundation/trust-type of organization, or they claim citizenship in a progressive jurisdiction with good relations to the less progressive one.

Many AIs are equipped with overrides to prevent certain actions or to enable legal authorities to take control over them. "Police mode AI" is a hotly debated concept on Luna, where many think this both infringes on AI autonomy and poses a risk to democratic society.

Advisor grains

Advisor grains are all the rage in many habitats. Small, artful aerostats connected to an AI that can give advice, act as a personal communicator and documenter.

Autassistants

Most orbital professionals use AI assistents to handle their complex lives. The autassistants are somewhere between a secretary and an extended mind.

Automated research

The problem is not that AIs cannot come up with research questions, it is that they are interested in so otherwordly things. Constraining them to come up with "useful" research tends to produce mere improvements on human questions. However, to a working scientist it is a boon to have an army of "igors" testing most variant approaches 24-7.

In some areas the concept of "emergent alphabets" are in vogue. All beings create models of the world, and these models tend to have a form of modularity – an alphabet. These alphabets of concepts and relations express what kind of world the being perceives and think within. Humans have one kind of alphabet, AIs have other kinds. By deliberately mapping out their emergent alphabet radical newperspectives on reality can be understood. Unfortunately, much of the emergent alphabets produced have an unclear relation to known reality. The result has been a growth of metascience, where researchers study the alphabets discovered by AI to find general laws or concepts that can be useful. Again, the results are usually not very impressive, but there are exceptions such as the Bulletproof Repair Platelet that have extensive applications.

Inspirations

http://ohhello.tv/index.php/work/view/microsoft_sustainability/
http://beyondthefold.net/Version_2.0.html
This is the information systems of the 2030's, 70 years ago. By now most places on Earth have something similar, often locally produced.

http://www.youtube.com/watch?v=ghzOKEegH80
Robotics