Humans and evolution have tried to solve many of the same problems---storing energy, turning sunlight to energy, moving physical joints, detecting photons, computing...

Humans usually solve these problems worse than evolution. But compared to my initial estimates, the gap between human engineering and biological artifacts seems to be surprisingly small and surprisingly consistent.

I think the typical pattern is for human artifacts to be 2-3 OOM less efficient than their biological analogs, measured by "How much energy/mass is needed to achieve a given level of performance?" Manufacturing costs differ by a further 2-4 OOM.

Ultimately it would be good to have a better sense of the distribution here (i.e. how common is "more than 3 OOM worse"), and to make a more serious stab at a "representative" sample (though it's hard to say what that means).

Performance comparisons

How Much Worse are Human Engineered Artifacts than Evolution		
	Performance (powers of 10)	Metric
Dialysis Machine	3 worse	Energy cost
Artificial Hearts	1-2 worse	Energy cost
Solar power	1 better	Efficiency
Solar power	4 worse	Payback period
Chemical Energy Storage	2 worse	Payback period
V100 GPU	1-2 worse	Flops/w
V100 GPU	4 worse	Manufacturing energy cost
Photodetector	3-4 worse	Overall performance given power
Locomotion	2 worse	Energy cost

Methodology for performance comparisons

Manufacturing costs

I think the energy cost of building organs is dominated by the cost of the fats and proteins from which they are built. Synthesizing the appropriate building blocks perhaps adds an extra 15-30% (discussion here).

When comparing to manufacturing costs for human artifacts, we could either compare the manufacturing costs to the 20% of energy required for the assembly of tissues, or compare the total cost (which is dominated by the raw ingredients for human tissues, and varies for artifacts).

I'm inclined to compare total costs to total costs when comparing the efficiency of downstream products (like solar cells, motors, *etc.*) and to compare manufacturing costs ex materials to manufacturing costs ex materials when comparing the efficiency of manufacturing processes themselves.

I don't feel great about this estimate for manufacturing costs of biological organs. But I feel relatively confident that it's an underestimate by less than 1OOM, since (at least in folklore) predators are able to recover about 10% of the energy consumed by their prey. So it seems like this could be at worst an underestimate of manufacturing costs by half an OOM or so on average over organs (though it could be significantly worse for some particularly expensive organs).

Embedded energy

Rather than computing the total energy embedded in a human artifact, I'm often just going to consider the cost and then convert to energy at a rate of $0.05 per kWh. The price of oil or electricity at the point of generation is about $0.03/kWh. Typically energy is a minority of the spending (and I will neglect all spending other than energy for the body), but this still seems like probably the fairest way to do a comparison. It might overstate the efficiency of evolution by up to an order of magnitude or so.

Artificial organs

Dialysis machines

Dialysis machines seem to use 2-3kW of power (here, but not very scientifically literate), and you need about 10h per week (here). That's ~6% utilization = 170W on average. The human body uses about 80W. I don't know how much kidneys use, but I think they are like 0.5% of mass and so a natural guess is about 0.5W. So dialysis machines use maybe 2-3 OOM more energy than a kidney to do the same work.

Comparing cost is harder. A dialysis machine can support 15 people, so let's call it 15 kidneys, and seems to cost something like $15,000, so about $1,000 per kidney (=20,000 kWH at $0.05/kWh). A kidney weighs about 150 grams. I don't know how much energy there is per gram of kidney, if I use the 1.5 calories per gram estimate from the liver this would give 225 calories ~ 0.2kWh. That's about 5 OOM cheaper than the dialysis machine.

Some substantial fraction of the dialysis machine cost is IP, licensing, *etc.*, and my best guess would be that the difference is closer to 4OOM than 5OOM. I'm making this adjustment here, and not in the case of other manufacturing costs below, because I have particularly little faith in the efficiency or cost-centric development of the dialysis machine industry (it has something of a reputation for being particularly inefficient even compared to the broader medical industry, which already has quite a bad reputation w.r.t. efficiency).

Artificial hearts

Artificial ventricles can be driven by a 6kg pump, vs. about 0.3kg for the whole heart, so maybe 20-50x more massive.

The heart consumes about 2.5 watts, about 2 watts from the left ventricle. Based on some googling I think that an artificial heart draws much less than 50 watts.

My overall expectation is that artificial hearts are within 1-2 OOM of hearts in terms of efficiency at runtime, but there are probably very large differences in manufacturing costs.

Artificial lungs

I expect artificial lungs currently in development to be within a few OOM of efficiency on most metrics at runtime, with relatively large differences in manufacturing costs. Here is some discussion.

History

I don't feel good about my ability to figure out how cost effective artificial organs were at different times. My impression is that dialysis is the easiest of these three organs, and became reasonably economical in the 60's. Hearts and lungs seem to have become possible later, and perhaps very recently.

Other organs

There are plenty of organs other than heart, lungs, and kidneys. Many of them are very tightly organized into the rest of the body, or have functions that are difficult to separate out in a modular way (or that make performance evaluation difficult). In many cases, a 2 OOM drop in efficiency would make a system completely unworkable as a medical prosthetic and so there isn't much investment and no one has tried to compete.

Here are some examples of other organs that it seems like are "fair game" for replicating:

- For the liver, it seems like many functions aren't characterized well enough to be plausibly replicated (though one might worry that this will tend to be true precisely for "hardest" things that biology does). The energy storage function of the liver is discussed under "chemical energy storage" below.
- For the intestines, it's hard to describe exactly what the function is. If we view it as analogous to turning biological material into electricity (including filtering out usable parts), then we can probably do relatively well. But I think building a small intestine that is a usable substitute within a human body seems a lot harder.

- The reproductive system seems particularly challenging, and probably is best to consider under the “manufacturing” heading below.
- The brain will be discussed under “computing” / “the analogy to intelligence” below.

Solar power

We can compare modern photovoltaics to photosynthesis in plants, on several dimensions:

- **Efficiency**: photovoltaics seem to be a little bit more efficient than photosynthesis. (Pop science discussion here.)
- **Density**: Rooftop solar panels weigh about 40 pounds for 60x156 square millimeters ~ 10 square meters. Rice leaf mass per area is about 20-50 grams per cubic meter, about 50x lighter.
- **Energy cost / payback period:** The most efficient water plants have doubling times (and presumably payback periods) around 1/2 day. Utility scale solar costs around $1 per watt of installed capacity (here). At $0.05/kWh, that’s a payback period of about 20,000 hours = 4000 days at 5h/day of peak-sunlight. This suggests that solar power payback is about 4 OOM worse than plants. Of course this comparison is complicated because the relative prices of different inputs are wildly different and so it may not be coherent to use energy as a common denominator.
- **Manufacturing cost.** If rice leaves are comparable to spinach, at about 0.2 calories per gram, then 10 square meters of leaves is 40 grams is about 8 calories ~ 0.01 kWh. 10 square meters of solar panels is about 2kW = $2000 = 40k kWH at $0.05/kWh. That’s about a 6 OOM gap.

That manufacturing cost calculation suggests that the payback period for plants should be about 1m, which is more than 2OOM lower than the actual fastest doubling times. I currently don’t have much confidence about my estimates of manufacturing costs.

History

Based on this informal blogpost I think solar panels became practical in the 60’s, with a price ~2OOM higher than today and a density maybe 1OOM lower.

Chemical energy storage

We are able to turn electricity into chemical energy at around 40% efficiency. Adding this on top of solar cell efficiency it still remains above the efficiency of photosynthesis. This is comparable to the efficiency of the (glucose -> ATP) transformation or (protein -> glucose) transformation. I don’t know how it compares to (glucose -> glucogen), (glucogen->triglycerides), or the reverse transformations

For the purpose of comparing energy costs, I think the best comparison is probably between (power to gas)+(gas liquefaction) and the (glucogenesis)+(fatty acid synthesis) in the liver. This comparison is still not a very close comparison, since one is turning glucose into triglycerides (via glucogen), and the other is turning electricity into liquefied methane (via hydrogen and gaseous methane), but they seem to be functionally analogous roles with similar levels of usefulness. I think this comparison unfairly disadvantages humans, since electricity can be more directly produced and used and so seems more analogous to ATP than to glucose.

Power to gas seems to cost something like $800 per kW of capacity (modeling assumption used here), which is about 15,000 kWh at $0.05/kWh. That means the system would have to run for about 15,000 hours to store as much energy as it took to create. For long-term storage with comparable efficiency to fat tissue in a human, you would need to liquefy the resulting natural gas. LNG plants cost about $500 per ton per year (here), and a ton of natural gas is about 15,000 kWH, which means you would need to run a plant for about 8 months = 6,000 hours to liquefy as much natural gas as the energy required to make the plant. That gives us around 20,000 hours for the system to "pay itself off" in this sense.

I don't know how to really estimate how efficiently the liver converts sugar into fat for storage. I could make a wild guess that 20% of the liver's machinery is usable for this purpose, and that the peak rate of storage is maybe 1000 kilocalories per day, that would be 5,000 kilocalories per day for a liver's worth of machinery. The liver weighs 1.5 kilograms. I don't know its composition, but animal livers seem to be 30% protein and <5% fat, if we use that estimate we get a bit less than 0.3 * 4 + 0.05 * 9 ~ 1.5 calories per gram of liver, which gives us 2,250 calories in the liver, call it 2,500 after manufacturing costs. That would mean a liver needs to run for about 12 hours to store as much energy as it cost to produce, about 3OOM cheaper than the human version.

Computing

A V100 delivers about 50 TFlops for 300W. A human brain delivers maybe 1000 TFlops for 20W~300x more efficient.

I think this probably overstates the efficiency difference by a factor of ~1OOM:

- A V100 seems to perform more communication per FLOP, and I think it will be possible to increase arithmetic intensity almost an order of magnitude while still retaining enough communication to simulate something like a brain.
- A V100 operates at 16 bit precision, while 8 bits is probably a better fit for the operations being performed by the brain. (This is another factor of 4)
- 1000 TFlops is an aggressive estimate for the brain, I'm uncertain between 100-1000.

My best guess is that all things considered computing in silicon is currently about 1.5 OOM less power efficient than the brain, and the other 1 OOM is due to structuring the computation differently (in a way that is easier to program).

Today the total cost of using a GPU is very large relative to the power required to run it. But a large part of that reflects R&D costs (maybe a factor of 3), and another part reflects the very short amortization schedules due to the current high rate of tech progress (maybe 1 year, rather than a more reasonable 5 year schedule). Both seem like they should be excluded.

Overall I think that 1-2 OOM seems like a reasonable estimate for the relative performance of computers and brains.

A brain weighs about 1.4 kg. Based on nutrition information for animal brains I'd guesstimate that a brain is about 10% protein and 10% fat, meaning that the energy cost is about 1.3 calories per gram. So that's about 2000 calories (~2 kWh) to build a brain. If we guesstimate a manufacturing cost of maybe $3000 for a V100, that's $60,000 for twenty voltas = 1.2 million kWH at $0.05/kWh, about 6 OOM of manufacturing cost difference. I think that's probably again an overstatement by an OOM, so the real difference is around 5OOM.

History

In 2000 a celeron processor delivered 500 MFlops for 20W, roughly 4 orders of magnitude worse. I think the real gap is more like 3 OOM over that interval, and probably 1 OOM per 5 years. I think the manufacturing cost gap is similar.

Photodetectors

Good resource: [Towards a digital camera to rival the human eye](), especially figure 16.

Possible photodetectors carve out a complicated space of tradeoffs. The most "balanced" photodetector in that paper is Image sensor 4, a 2003 CCD camera from Atmel, which uses about 1OOM more power, 1 OOM lower temporal resolution, and 2 OOM lower precision than the human eye.

Intuitively we'd want to take a product between these---most of the properties trade off linearly against power consumption. If you take a product, you find that most cameras are about 6 OOM below the human eye, and the Atmel is 4-5 OOM.

However, I think that's a very dubious methodology given the way they measure the performance of the eye. For example:

- The dark limit of the human eye is calculated as 0.001 cd / m^2, because below that brightness the eye can no longer recognize colors. But it seems quite likely that before the eye loses the ability to recognize colors, it begins to drop spatial and temporal resolution.
- The resolution of the eye is reported at the center of the visual field, but the size of the visual field mostly consists of areas where the resolution is much lower.

- Similarly, evaluations of temporal resolution are based on ability to detect low-resolution but high-frequency patterns. If different cones fire at different times, then we would have lower spatial resolutions at higher temporal resolutions.

So I think that multiplication is a little bit too generous to the human eye, and the Atmel is probably 3-4OOM worse than the human eye all things considered (ignoring manufacturing cost). I don't know how much progress there has been over the last 15 years; apparently not that much over the 5 years from 2003 to 2008 since none of the cameras from that period decisively beat the Atmel.

My best guess is that photodetectors are about 3 OOM worse than the human eye, depending on how you count, ignoring manufacturing cost.

A modern cell phone camera costs maybe $15 (=600 kWh at $0.05/kWh); it has about 1 OOM fewer pixels (10M vs 100M) and comparable temporal resolution. I think an eye contains a few grams of tissue, so is maybe 5 calories = 0.005 kWH. So the cost is about 6 OOM lower after controlling for pixel count.

History

I think analog cameras perform *much* worse than the eye, typically having ~3OOM fewer pixels and having either very low temporal resolution or very high power consumption. Digital cameras were only seriously introduced in the 80's (discussion here). A 1 megapixel camera in 1991 cost $10,000, about 4OOM worse on price/performance. Raw performance by the early 90's was maybe only 1-2OOM lower than modern cameras.

Actuators

Would like to look at:

- Motor speed/strength as a function of energy/mass.
- Efficiency of flight
    - Smallest drone: 40 millimeters across (here)
    - Smallest flying insect: .15 millimeters across (here), 300x smaller
    - (There are better metrics)
- Robot arms
    - UR5: weighs 20kg (but power source is external), lifts 5kg, 200W here
    - Human arm: weighs 4kg, lifts 20kg, ~20% efficiency
    - (There are better metrics)

Manufacturing

Materials

Manufacturing costs in the body seem to be on the order of ~1 calorie per gram = 1 kWh per kilogram.

This table describes the total energy cost for a range of materials (1MJ = 0.22 kWH). It ranges from around 1 kWh/kg for simple materials up to ~20 kWh/kg for synthetic materials like PVC or aluminum.

Human tissues are typically much more complex than raw materials, but most of the energy seems to be synthesizing low-level structure rather than assembling complex high-level structure.

Complex structures

As far as I can tell, animal bodies are extremely efficient at constructing complex structures at the scale of the cell and larger, but this is really expensive for humans except in special domains where we have been able to drive costs down. Ratios of manufacturing costs above are already quite large (typically ~6OOM), and probably they would be larger if not for the resource costs of the low-level structures.

To estimate ratios of manufacturing costs proper, it would be good to consider the setting where the human body is provided all of the necessary raw ingredients (especially fats and amino acids) and does not need to synthesize them. I expect this would typically add about 1OOM, with costs now dominated by the cost of synthesizing proteins from amino acids.