An Introduction to Live Interop

Live Interop

Beyond Space

Space and modularity

- When building with LEGO, you use a collection of bricks, each serving a particular purpose—a small brick to create details, a flat piece to form a base.
- These bricks can be arranged into complex structures to create a model. In a sense, these bricks and steps are modular components of a composable system.
- The modularity of these bricks arises from their physical instantiation in space. Any given brick can be picked up and used to build a wall after it finishes forming part of a car.
- Bricks don't proactively assemble themselves – they're not like the enchanted brooms in Fantasia, automatically seeking out the next spot to be placed – you have to do the hard work of figuring out which bricks to use and when.

The skeuomorphism of space is overreaching.

- Skeuomorphism (the design aesthetic where an abstraction retains elements of its originating environment) has significantly influenced the design of our digital environments. Many logos, fore example, have retained elements of physical space. However, the influence of skeuomorphism isn't only superficially evident in icons; it has bled into the core of how we engage with our digital tools. Features intrinsic to the physicality of space, like its modularity and composability, have been reconstructed as features of the non-physical digital realm.
- Take Unix tools[1]. Like LEGO bricks, Unix tools are modular and composable. They do one job and do it well. They are intended to be composed into more complex tasks: use one tool to read the file, another to create a copy, and another to move the copy to a different directory. They don't try to think for you.
- Or consider Word Processing. This text was drafted on a Google Doc - this is where the text "was". The Google Doc was a space, in that it had a lot of local scoping. The only people who could read it were myself and named individuals with whom it was shared. The text wasn't accessible when I opened my Signal app. I had to go back 'there' (to the Google Doc) in order to reread it - I had to digitally walk to the space 'where' I kept my data. I can add a new channel, and put some different text there.

[1] Unix tools are a collection of small, focused programs widely used in operating systems like Unix, Linux, and macOS. Each tool has a specific purpose, making them composable—the output of one tool can be the input of another. This allows for building complex workflows by chaining them together. Some examples of Unix

The tyranny of folders

- We started with folders, and I suppose are still mostly stuck with folders. When you put a file in a folder, that's its location. This was a genuine constraint from physical folders and lockers. You can't put the same file in multiple folders! To retrieve a document, you have to go to the single physical space that you kept it at. Functions define partitions on the domain indexed by the image, which is a skeuomorphism of literal space partitions.
- Yet I can't imagine organising my Roam notes without tags. Notice how tags are not a function: File -> Folder, but a relation File -/-> 'Folder'/Tag. The difference is that every item can get zero or more labels in the tag model, whereas in the folder model every file must be in exactly one folder.
- The single-valuedness of a function is where 'space' emerges from - you put things in a slot. Fermi-Dirac applies, not Bose-Einstein. When filing documents, each item has one specific place. It's similar to how in quantum mechanics, fermions (like electrons) cannot occupy the same state, as opposed to bosons, which can.
- I claim this is actually a want for modular determination bias. That is to say, you want one answer to 'which folder should I put this in?' because limited/local/urgent things want to keep things simple, determined, workable. This suggests a preference for straightforward, unambiguous categorization systems in situations where quick decisions are necessary, akin to needing a clear organisational system during a time-sensitive project.
- This seems to be motivated by a demand for grounding. The question you ask when you want a file is where is the file 'grounded'; what is its 'true' location.

Liver interop breaks down artificially imposed constraints on storage

- When we consider the evolution of information technology, the perpetuation of space-based design aesthetics isn't surprising. [example]. [The writer is pinched into space metaphors to pre-emptively do the retrieving].
- But, as already stated, these aesthetics reconstruct artificial versions of physical-world constraints. And as technology develops, the [opportunity cost] of these artificial constraints increases. LEGO is not going to be imbued with Fantasia-esque [liveness]any time soon; but this story doesn't look too outlandish when considering what programming will look like with AI copilots.
- [promise of liver interop] And in so doing, hopefully ensure that our relationship with the interop is more constructive than Mickey Mouse's relationship with his broom.
- Live interop will work with a single 'store' of raw data and will automatically generate indexing markers (e.g., tags). This will enable the retriever (the user seeking to access the information) to dynamically pull up a context-specific dashboard, such as a 'conversations with boss' dashboard. The originator (the user who deposited the information) doesn't have to categorize information based on the anticipated needs of the retriever; live interop empowers the retriever to do this work themselves. The burden of contextualizing information to suit the retriever's needs is much more shared.
- High-actuation storage, where copying is cheap, or where retrieval can be a low-latency 'smarter'/'more alive' worker than just 'dead and dumb' physical space, allows

implementation of the 'magical' idea of relations. Instead of going to the space, you ask a system to (intelligently) retrieve it for you.

Privacy is of course something that we should be able to each decide, like we already do when we make a note in a 'private space' vs a 'public space'.

- Digitally, this is not usually spatially separated, but cryptographically. Even airgapped hardware wallets are dependent on the cryptography working out. For any question or curiosity you have, there is a lot of potentially relevant data out there. The salience of that data will often not be determined by 'actual' relevance (the telic constraint), but instead by the extent to which the 'spaces' where the data is stored can mix together (causal constraint).
- [examples of organisational silos]
- [Remember this being the start of high-actuation spaces? It was in a wildest-speculations thread, about the causal constraints becoming secondary as construction becomes cheap.]

This is all about cycling attention/salience (exobrains style).

- Spaces and Things are epiphenomena; our division of reality into defined spaces objects is a mental construct - the universe has no rule distinguishing between a chair and a stool. But the fact that these concepts are constructed doesn't make them meaningless; we have a need for solid anchors (determination bias). Regardless of your views on whether the *Theseus* is a different ship after it had a plank replaced, you still need to sail some ship to your destination and you're probably going to use the same name to refer to it throughout the voyage in any case. Even if the storage 'space' for your research material into canine evolution begins to lose its fixedness, you might still want to it all in a brown folder with a furry stamp on it.
- But once we move away from relying on [artefacts] inherited from the physical world to organise our information, our existing ideas of space [can be FGF'd + constructed.] We can begin to think about the overall market for retrieval instead of focussing solely on the supply-side architecture of storage. Retrieval (hitherto "space") becomes a domain resembling an economy of (habits of) attention, not solely a domain of (supply side) packing problems.
- The pipeline for "storage" right now is as follows: originator investigates where to [store/deposit] information for anticipated retrieval → originator [stores/deposits] the information; this needs to change to: originator [stores/deposits] the information → originator leave [hints/markers] for the retriever (if they like). With heavy live interop, we don't even need to have modular 'space' for text that's yours vs text that's mine. Already, our centralised storage (like Google drive) doesn't really separate those. But it still isn't 'live' enough; I should be able to 'recall' your thoughts, translated to me in my jargon.
- By freeing our engagement with information from the constraints of modular space, live interop will enable us to begin to leverage our deeper subtler attentional recycling patterns to enhance the retrieval process. The brown folder with a furry stamp doesn't exist as a fixed segment of modular space; it's created as an [aesthetic/telic device] to situate and sharpen your attention. [Live interop will dynamically tailor…]

In this case, we're talking about reconstruction. Stochastic parrots already merely retrieve information, with some recombination abilities.

- Ted Chiang famously preferred "blurring" instead of "recombining". But that doesn't matter much. What is useful is his lossy vs lossless picture.
- Here I'm talking about intelligent lossless retrieval. This isn't some holy grail of Aligned AI; It's not solving some general problem of deception in advanced intelligence. It's only allowing retrieval.
- In fact, soon it will probably become Soryu's AI4 (or was it 5?) and will become sticky exothoughts that you have to let go of somehow. It's the simplest thing you have ever done (just GET OUT OF THE EXOBRAIN) but surprisingly hard as it seems to be to quit scrolling Twitter.

The profunctor revolution (or "relator" revolution) is all about turning functions into relations

- To switch to tags instead of folders is, as a reminder, about having a parent relation instead of a parent function, for each data node. In practical terms, this is like using hashtags to organise content on social media instead of putting posts into a single folder. A post can have multiple tags and be found through various connections.
- The profunctor revolution, or 'relator' revolution, is all about turning functions into relations (and also stepping up on the categorical periodic table). This means that instead of viewing processes as one-way functions from input to output, they are seen as two-way relationships that can be reversed.
- Global retrieval means working more holistically, working with a lot more context, the bane of 'interpretability' (at least when not in coexplication style). Skill in connotations is essential, you need to sense a vibe of a room to see if it's coming together. This implies that understanding the big picture requires considering all aspects and nuances of a situation, much like a detective piecing together clues from various sources to solve a case.
- Conflict resolution is also perturbed by urgency that wants to determine "I am so tired of this, what is that you want out of this? Do you even know?" without realizing that coherence is a property of polices, not preferences. And mixing that up in attempting to urgently get to a shared policy instead of spending a first stage on patiently legibilizing contradictory preferences. Ironically then, this is campaign to become more spacious (in the sense we've used before) than the restrictedness of "space" (in the sense that the profunctor rails against).
- Ironically then, this is a campaign to become more spacious (in the sense we've used before) than the restrictiveness of 'space' (in the sense that the profunctor rails against). In essence, this is advocating for a more flexible and expansive way of thinking and organising, beyond the limitations of traditional, rigid structures, much like preferring an open-plan office to cubicles for better collaboration and flow of ideas.

Beyond Text

Reading as a dynamic process where the local context of the reader is enmeshed with the originating text

- Currently, when we want to read content online, we are forced to accept the text in the context presented to us by the originator. Live Interop will give readers control over this experience.
- Rather than having information fixed and pre-determined, Live Interop will enable it to be dynamically adjusted to suit the reader's local context. This is largely within the capabilities of current LLMs, and there may already be an app that provides this functionality.
- Initially, this could take the form of a website or browser extension that allows users to simultaneously redraft a target text into a summary, light version, stern version, math-free version, etc. Users could pin their favourite text-filters on the dashboard with a drag-and-drop feature.

Writing as 'recipe' formulation,

- where the author outlines their originating ideas alongside the construction process.
- A writer might begin with a sketch of bullet points. They could then select these bullets and click "flow," select a paragraph and click "light," then select a bullet and click "emphasize." The application would retain the originating 'recipe': the text and the series of transformations as chosen by the user (the flow → light → emphasize steps). Writers will be able to select a paragraph or text and adjust the funniness, flow, sternness, etc.
- This approach would allow readers to receive the 'recipe' and have it adjusted based on their personal preferences. For example, emphasis might be highlighted in a way that works for them, like triple bold for someone with ADHD, or a slightly stern tone for a British reader. More live interop and exobrains are probably relevant here.

The craft of writing will be decoupled from the art of writing.

- Once we have a bunch of satisfying "recipe - text" pairs from consumers, we can train an AI to infer recipes from existing text, thereby making live interop even easier. We could also modify recipes, *open source text filtering*, so that we can remove "woke euphemisms" or "lazy jokes" or whatever or other undesired elements.
- These recipes will also double as prompts for LLMs, allowing both consumers and various LLMs to have their specific characteristics worked out by "low-level engineers" and modified accordingly before being fed into an LLM.

Once this is underway, the socratic style teaching of technical fields will be much closer.

- For instance, you might be given problems like Bongard problems (e.g., convex spaces on the right, concave spaces on the left) and asked to name them in a way that makes sense to you (like "kiki/bouba").
- You will forever use those terms, and every textbook you read will use those terms. No more double penalty of memory and insight in learning. No more "have to" of learning any standards. Standards are 'dead interop'! Boo!
- Although I'm also already feeling the profaneness from automation. It's like "GPT could have made that joke, ugh" or "bah, this sentence has a GPT vibe" I suppose there might

always remain some implicit "labor theory of value" in art and human meaningfulness, that will keep the goalposts and eyes rolling forever far into the future.

Beyond Time

Live interop will free information storage from temporal constraints

- Imagine you join a meeting on Zoom a little late. AI summarizers can "catch you up". Instead of immediately chatting, you peruse the transcript/read the summary.
- Now that's great, but a little boring. Ideally you'd have an "immersive summary", a mini-meeting that you get to see live via a generated video summary. You watch that, and then seamlessly transfer to the "real" meeting.
- Except when Alice references a detail that wasn't given enough attention in the virtual-summary-meeting you initially watched to catch up. So, before you get confused as Alice starts using the details assuming you were in the meeting, the interop again pulls you back into a simulation where Alice says "so, just to recall, some details I want to add…" before dropping you back into reality.
- And in fact, your replica will already have joined the meeting even though you were 10 minutes late. And you watch it, and might say "oh, no, that's not what I would have said!". And so, once you say that, *everyone else* will be put in a simulated meeting where your replica says "I know I said earlier that X, but no, I take that back. What I actually meant to say was ...". So everyone is again caught up.

Rather than being blocky or jarring, imagine a perfect stabiliser. Like in a drone or in video-correction.

- It's uncanny, that despite all the noises, rapid responses can correct destabilisation quite magically and beautifully into a coherence. This blurs "construction" and "reality", doing such a fantastically better job that you will only occasionally need to rebel against this matrix.
- The real-constructed dichotomy will make less and less sense; not in a way that is magical, merely apparently magical (for someone to whom the technology is entirely novel). For this experience only the following criteria needs to be fulfilled: a) user experience is seamlessly stitched, and b) temporal simultaneity is optional. (Notice, of course, that it will be considered "unusually embodied" for you to be breaking out of such a Matrix of connection FGF in everything)

Our relationship with time will change to be more fluid and less constrained

- Right now, our ideas of what constitutes the ‘past’ are quite rigid - even more rigid than a scientific account might suggest. We treat time as a fixed (causal) constraint arising from the fabric of the universe. When someone first encounters special relativity they are often shocked that the universe has such a flexible approach to time.
- Live interop will mean that our relationship with the past will [change] to become more fluid. Rather than the past constituting a fixed element of the universe (a causal constraint), it will increasingly become an idea constructed to serve our own purposes (a telic constraint). In the same way that a film director might intentionally stretch and contract time as an artistic device, we will be able to do the same as an attentional device.
- There will still be a "base" reality of time, but we won't really care. Just like we don't really care about individual quark changes except rarely for some engineering. Most of our currencies operate within a much more human-reality... which will also fluctuate .*Edit:*

Matt suggests "constructed-yet-real-time" to replace "real-time" in a world with more live/high-actuation orientation to time.

Beyond Isolation

The future of art and media is one that is more flexible, dynamic and unconstrained

- Most people have already seen a dramatic change in the nature of our relationship with art and media in the scope of their lifetimes.
    - Not so long ago, watching a movie required physical attendance at the cinema. You could choose from a list of 10 films if you were lucky, and had to attend a prescribed place at a prescribed time. Too bad if you happened to miss the showing, or if the film wasn't released in your country.
    -
    - And as mobile phones and high-speed internet have become increasingly ubiquitous, Netflix and streaming platforms have meant media consumption has been - for all practical purposes - freed from the tethers of space and time. The primary limitation we face is the content that is available.
- Live interop hopes continue this trend, until the primary constraints on the media we consume reflect the telic constraints of the artistic genre rather than the causal constraints of its method of production and dissemination./
    - Soon, rather than sitting down and sifting through the litany of algorithmic recommendations to find a suitable film, we will be able to convey
    - Maybe you ask the interop to create a new season of Narcos to watch tonight, but have it take place in Syria with Brad Pitt, Mr. Beast and Travis Kelce in the leading roles. Or maybe you have it

To some this might feel strange. Will people actually have the motivation and capacity to create their own media (even with AI assistance)? And if they do, won't this lead to the breakdown of shared culture?

- Exercising imagination isn't always easy - it requires focus and energy. Plus, it atrophies without use.
    - Until we can read minds, maybe most people are going to be too boring and/or lazy to actually bother to create media. Instead, perhaps they'll come home and want to be spoon fed a show/movie/music. Value accrual would happen at the distribution end (Netflix,Spotify, etc), since they already know your preferences.
- Plus, even if we do have the requisite capacity for creativity, wouldn't this then just lead to a breakdown of common culture?
    - If every piece of art and media is tailored to the individual what does this mean for common standards of truth? How can you talk with anyone about a film, or a meme, or a song if everyone has their own tailored

These fears lack imagination. Creativity will be facilitated, and culture shared, at the level of inputs rather than outputs.

- Imagine saying "you can't talk to other people because they have totally different experiences" - this is nonsensical.

    - What it means to share in culture will change; just like live theory, we will have live theorising. We will exchange recipes rather than outputs of conversations, and thereby be able to have an AI make the same flavour of point that I want to for the movie I saw, translated to the movie you saw.
- And this will help solve the challenge of creativity. We don't have to generate ideas from scratch, but can immerse ourselves in a community
    - When someone learns to cook, they start out with simple recipes that they rigidly adhere to. Over time, out of necessity or curiosity, they can begin to pay less and less attention to the recipe and start to experiment with their own substitutes or even develop their own version of the recipe.
    - In the same way, we will be able to hone our creative capacity by leveraging the work of others. Maybe you detest shaky cam, and so start by altering films to suit your videographic preferences. Then you add in components from the book that were cut out for timeliness. Then you shift the plot into a different historical era, or make it into a musical, or have it incorporate your likeness as an actor.

And new forms of creative expression

- So many casual constraints limit our ability to …skill,
    - Imagine learning to play saxophone, and having your attempts at improvising backed up by …
    - Or seeing you child's short story brought to life
- And experiences no open to anyone
    - What it's like to conduct an orchestra
    - Fantasy world-building game
    - What would it be like to be the consul of th roman republic. Could you avert the collapse anbd emergence of th empire

Consider how music has evolved

- Owning a few vinyls of well-known songs to play at home -> keeping a mini-library of CDs for artists you like in your car -> Maintaining a large database of specifically chosen songs on itunes -> Maintaining access to practically every song ever written on Spotify
- Arguably we've seen the fracturing of musical popularity into many smaller independent artists. But this hasn't caused a fracturing of society
- And when we want to have shared experiences, we go to concerts together - hence the surge in the popularity of live events.

Beyond Asymmetry [unedited unformatted]

- ***The Currency Problem***
    - In "Guided by the Beauty of Our Weapons," Scott Alexander how logical debate serves as an asymmetric weapon - one that gives advantage to those pursuing truth over falsehood. But what happens when our technological capabilities make everything equally compelling, regardless of truth value?
    - That which is sexy has currency - it captures attention and spreads. The UI fantasy here is to create a translation/exchange system that could convert

low-volume, obscure mathematical adventures into high-volume, rich-CGI real-world adventures. In other words, we're seeking more integration through live translation (currency exchange).
    - This presents a fundamental challenge: When we can make anything vivid, visceral, and awesome in high-res embodied CGI without regard to truth (i.e., symmetric weapons), how do we maintain asymmetric advantages for truth-seeking?

- ***The Licensing Challenge***
    - I'm curious about when symmetric weaponization will be ritualistically licensed as asymmetric, with intelligent/live/connotative justification in the licensing rather than "straightforward" deductive proofs. We might see social infrastructure become as opaque-yet-functional to end-users as most devices and software we use today.
    - This is already happening in primitive forms - generative AI is trained to avoid certain topics and images. But I'm thinking of something more sophisticated: a licensing process that is more "automated." By "automated," I mean "intelligent" and "live interoperable" - terms that somehow feel like the opposite of automation.
    - The thing already self-censorsLike an automated appeal process to get the LLM to do something it is currently refusing?"I'm only generating images of my girlfriend, this should be fine"

- ***The Full-Bodied Future***
    - Imagine being immersed in VR or jacked into the Matrix. We can use interfaces that your "less attentive" parts of your brain are used to (like Twitter) to inject useful and wise information (like Calculus). This represents a major decoupling of interest and information.
    - Traditionally, your aesthetic judgment works fairly well to discern information. But when we can make anything interesting, we face both opportunity and threat:
        - - Opportunity: You might finally learn calculus, which doesn't look shiny but actually is valuable
        - - Threat: We might make Time Cube look equally compelling

- ***The Intelligence Layer***
    - You might want to make up filters that will decide "legitimate" and "illegitimate" motivation, now that we can just "pay" your hindbrain in its currency, to get its attention.
    - But this is too dead, old ideas of interop and filtering. Uses solely cryptographic certificates and mathematical proofs or a small canon of ethical things that humans can verify.
    - In a nore full-bodied version, every bit of viscera entering your body will have live intelligent filtering - essentially an AlphaZero version of legal, tailormade for you.

But this isn't just AI - calibrated reputation networks from human communities will be crucial inputs.
    -
    - The "tailormade" aspect is dual-use technology. The nightmare scenario isn't universal death - it's everyone trapped in their own personalized hell, similar to what some describe during bad psychedelic trips. It’s al ot scarier than us all dying together

- ***Taking Intelligence Seriously***
    - We need to think more carefully about "taking intelligence seriously" in high actuation spaces. There are two ways we typically fail:
        - 1. Limited imagination (conservative thinking)
        - 2. Unlimited imagination (relegating capabilities to the platonic realm and thereby not thinking about its reality)
    - The second failure is particularly subtle - treating extremely cheap as equivalent to free, for example,by assuming that currency will not be made out of zero substrate. Merely cheap should not be approximated to nothingness when dealing with dynamics of similar magnitude. Your epsilons might cancel out and give you finite values.
- ***The Infrastructure Challenge***
    - Current attempts at interoperability (like XMTP and Matrix's Element) struggle with basic chat integration. Maintaining interop across diverse systems feels pointless - how can one person handle the ever-widening conflicts of philosophy between various apps and projects?
    - An LLM-based slightly "agenty" coder could manage this tedious work. Your wish for seamless integration could be fulfilled: use your favorite UI, with live views generated from live backend interfacing, your notes collated in real-time, pulling up all relevant definitions in a "single app."
    - This represents extreme "substrate flexibility," where traditionally tedious barriers are "intelligently" overcome. But it also raises the question: when everything can be made compelling, how do we ensure truth maintains its asymmetric advantage?

Beyond Asymmetry

Beyond Storage [partially edited]

What is the evolutionary advantage of pain over reflexes?

But there is a residual puzzle: why not use reflexes instead? The patellar reflex allows the organism to move rapidly and effectively, but no pain is involved; same for the blink reflex. So, why must a stubbed or squashed toe be accompanied by intense pain—why not arrange the nervous system so that the foot is quickly and reflexively withdrawn but without the intervening agony? The pain sensation doesn't seem necessary to the function; it seems like a gratuitous (indeed sadistic) add-on. The only thing I can think of is that the pain is somehow necessary for ongoing flexible voluntary behavior in the presence of harmful stimuli, as in managing a broken bone or a burn. But though that seems true as a matter of empirical fact, it is difficult to see why it has to be true. Presumably it has just turned out over the course of evolutionary time that pain is a more efficient way of handling injury than a purely reflexive and pain-free method; but why this should be remains obscure. So, the existence of pain is something of an evolutionary puzzle, especially given its functional downside (its phenomenological downside counts for nothing in the evolutionary game). It clearly evolved over millions of years and is close to universal, but it's puzzling why it exists at all as an adaptation

https://www.colinmcginn.net/evolution-of-pain/Colin McGinn Evolution of Pain - Colin McGinnEvolution of Pain

The 'ongoing flexible voluntary behaviour' that pain (potentially) enables has a similar flavour as the ongoingness (we contend) is inherent in model-free

- Thinking you can get away without doing your maths textbook exercises because you've read through the definitions and remember them, is also a case of having a more [coarse-grained functional] model of yourself than is actually the case.
- The act of doing the exercises allows us to appropriately cycle our attention to lend the relevant maths concepts greater salience. There may be a whole cluster of proof techniques and properties encapsulated in a particular definition that need to be unwrapped and recalled at the right moment, to get a sense of it.
- You get embodied practice of propositions, which is easy to mix up, but is true. This is not just logical omniscience. This is also about habits, about which things will care to jump to mind.
- Salience is more on the telic side of things than the causal side of things.

Knowledge requires both description and prescription. Information is made prescriptive through its salience.

- You can have all the description, but without prescription (which is normally assumed obvious, and therefore hidden) it is pointless. To have pure description is to merely have a textbook on your table that you could read, but never do. That's a kind of knowledge, this access to the textbook.
- But to have integrated it is for the results to speak up automatically. To "want" to jump up, have the power to steal your attention. Like focusing on a heartache steams up.

Information can also be made salient through augmentation, or even through a consultant.

- If your employer hires an expert consultant to aid you with a complex project, in some senses their advice is like an unconscious whispering of a relevant insight. But instead of it coming from within, it is coming from a human next to you.
- But the human consultant won't be as closely coupled with your currencies as your unconscious insight. You'll have to articulate your thoughts using speech or writing - forms of low-actuation interop - in order for the human to give you what you want. You'll have to "translate" out, into their currencies. And even then they may not *get it*; maybe they don't fit into the organisational culture, maybe they don't understand the institutional and political context. Finding a consultant who can *tune in* to unspoken organisational dynamics can be exceedingly difficult, but also incredibly impactful.
- Nevertheless, the process of "translating" your thoughts for another human is not quite as crude as "translating" them to a machine, where you have to physically/manually rearrange its substrate to get it to say something relevant. You can interact with a human in a "deescalated" manner, just translating the spirit of your intent (you don't have to manipulate their brain to communicate your thoughts into theirs). Your currencies of value are much more closely integrated, with much less friction.

Takeaway? - Live interop will present another path to salience; like a consultant who understands what knowledge you need and designs your information environment in order to leapfrog the need to seek prescription through instantiation in your subconscious. Means our capacity to interact with new information expanded;

Beyond Truth

- ***What does truth have to do with space?***
    - Truth and space both have a staticness to them, to the extent that truth is about invariants and deep underlying principles and that which is indisputable. Space is the thing that is orthogonal to time i.e. change. Truth, or knowledge, tends to be treated as that which you can *store*. Usually permanently, because it is simply true
    - What do you fill space with? Presumably, truths. Convergence of things held in a truthy way feels related to a "point" in space. "Ultimate" or "really" or "underlying" have a sense of "bottom", on which "higher levels" sit, which is a spatial resolving.The point I want to make is that we may have to give up on the implicit emphases on truth and space, *simultaneously*.
    - Above and elsewhere I've talked about 'referential reach': I propose that we no longer have 'dead and dumb' retrieval of 'Truth' any more than our files! Talk about wholes without having to talk about parts.Better q to ask: what is meaningful and useful given your current context?
    - What replaces "truth"? This is why *constructed-yet-real* in the high-actuation frame

- ***Hallucination-Compatibility and Nonmodular Integration of Truth***
    - This isn't quite live interop:https://news.ycombinator.com/item?id=39854921 But I expect stuff like this to have a development lifecycle/trajectory of setting up access to ports for a copilot, letting it work out how to dointerop well, and occasionally giving it advice on what to change (via natural language or code).
    - Instead of compile-time, you might have integration-time for interop-ing, and thereafter automated maintenance as the software evolves. Feeds of change logs will also be connected, and with temporary downtime in reintegrating the changes (but at the user-end looks like just latency of interop, since it's easy to cache user actions and fill them in).
    - And of course, similar to the #"more live interop" thread (on time and zoom), you can have simulated/hallucinated interactions in the meanwhile to avoid latency, with "eventual consistency" of your information.
    - Information doesn't need to be grounded in "reality" immediately, like your brain putting together various signals. Our interactions are mediated with a temporal mediator. This is a breaking free from temporal modularity of information
    - We can handle "hallucinations" as long as we're more patient, less myopic, more able to deal with changing positions rather than needing answers to be fixed.
- ***That "live eventual consistency" and "holistic/non-factored" have something to do with each other is not surprising.***
    - Deacon talks about the indecomposability of life, as opposed to machines. When the whole doesn't need to be a sum of its parts, it's freer to be other things.
    - Changing answers can be pretty frustrating. If I say things are X, then say, they're not-X, then change back, that's not nice to you. If "I want to eat Italian" changes to "Nah, I want to eat Asian now" after we drive 10 miles to the fancy Italian place, people will rightly sigh and be dismissive at you.
    - **But this is only in a low actuation regime.** You don't sigh when someone says "nah, I want a different desktop wallpaper." You just switch it, because construction is cheap.As we automate our followthroughs, knowing "oh we're using discord" was a hallucination and the live-interoper, on receiving the "real" information changes it to "sorry, that was an accident, we're on slack", you might be frustrated... unless you have your own bot that happily weaves between the two!
- ***This mutual buck-passing of frustration from stickiness, the tyranny of low-actuation, causal constraints, could fall away quite a bit!***
    - When everyone is dynamic, there is no frustration of "but I built my entire [blah] based on the decision [blah]!" because both [blahs] can be transformed into the new "decision" *automatically*. Not only your own frustration, but also the frustration of whoever you were supposed to deliver to, is dissolved (or minimized at least, like changing a wallpaper).
    - This is high actuation "truth". No piecewise reductionist assembly of truth. Just eventual integration. Patience, trust. Everything we do is interfacey (such as me writing this), and when they all have capacity to be fluid/ high-actuation/cheap-construction, when they're all live interfaces, when nothing

is like <#1195404287887552563> we won't have to worry about hallucination so much.
    - Your "it is sampling though!" is right. It doesn't need to be grounded. As long as it remains sensitive to changing "real world" and keeps moving us as close as we could have been if we'd received "truth", it's doing a good job. [This is Groundless Sensitivity, but for truth.]
- What will still need to be figure out, if we embrace "hallucinations" while working on their minimization: prediction of "the real world" and credit assignment (this is a big reason for keeping stores of truth)
    - Martin pointed out that the latter is dual to prediction. So explication is useful for prediction, and coexplication is useful for credit assignment(?) We will probably need these systems to also be hallucination-compatible. I think this is possible, if we notice how truthy they are, rather than dismiss them as "being untrue".
    - A lot of this is also about what generalization/extrapolation/integration means, and how we should deal with misgeneralization. Rather than condemn it and be fragile about it, the invitation is to use what's useful about it, and leave it correctable.This is likely a very big part of <#1196533078718435328> [Post-rigor]. It might feel chaotic and disorienting, but I'll invite you to see how we're already in post-rigor.
- Rigor is fragile.
    - Ex falso quodlibet is a logical principle. Make one wrong step and you can prove anything. Connotation is more robust in that way than denotation. So is post-rigor.
    - High actuation can be patient and accommodating with "untruths" as long as they are eventually integrated well.
    - Something generalizes in an ML model and something else doesn't... that means we have still something to impart! We're not surprised or upset when learning how to spell in Mandarin doesn't generalize to balancing on a beam. We're upset when the bot doesn't do the thing we *expected* it to do, dumb bot! IOW, we have something to teach
    - Let's build systems *to allow us to teach*, importantly, *teach patiently*, so that we get eventual coherence.
- Also related: the decentralization of modeling. Where you don't learn science and Facts from standard textbooks in school anymore than you learn software engineering from "standard software".
    - Those are all models, and we end up learning how to understand model-making as if there was only one, or one linear progression. It's good to have excellent example, not good to revere them as truth.
    - There is a doc [this one] where I attempt to explicitly connect the profunctor revolution's irritation with "space" and its irritation with "truth":
    - *Also related: truth debt/ forwardy debt [link]***beyond modular credit** *[Attenuated casual constraints of paywalls and access, as retrieval gets smart, with seamless microcredit. Debt is eventual integrity of credit assignment. Similar to the*

frustration of stickiness above falling away circularly, so might this. Fine-grained finance is obviously relevant.]