

What does it mean to 'generate' information?



[work in progress]

Polish-Bulgarian Workshop
PAS Warsaw, September 2023

Some motivations

- Information-talk everywhere in CogSci. But not enough focus on the *operations* on informational states.
- Neural information-processing lingo: selection, integration, amplification, attenuation... **generation** [some exceptions: Fazekas & Nanay 2021; Fazekas & Overgaard 2018]
- Old philosophical chestnut - how's new knowledge possible? (Plato et al.)
- And knowledge is made of information... right?
- The goal: provide a definition of what is information generation for a cognitive (esp. neural) system to contribute to this chestnut.

Some background

- Inference as ‘filling-in’ the gaps. Think inferential statistics (vs. descriptive)
 - In perception: ecological vs inferential [active, constructive]; filter-out vs construct from the gist
 - ‘Passive’ theories: all’s ‘in the signal’. To perceive: just make sense of the bombardment!
 - Internal representation *detached* from input (visualization, counterfactuals, internal models, recall etc.). They have to contribute as a source of information!
-
- Generative AI; Rumours has it: hot and new capacity. But what exactly is this generative component really? [following Buckner’s 2018 *modus operandi*]
 - Some other debates: e.g. consciousness studies (Kanai et al. 2019); concepts (Quilty-Dunn 2020)

Maybe it's already settled?

- Loss of information is generally understood in algorithmic information theory terms, i.e. defined over real patterns.
- But what about the informational *gain*? Lossy operations are not reversible!
- Gain is usually equalled with mutual information or KL-divergence. (Cover & Thomas 2006)
- But that's not generation. Neither it satisfies us in philosophical explanation what does it mean to gain new information.

Pump it (intuitions) up

- Intuition pump: to generate is to (at least) create new information.
- But clearly just churning new strings that merely eat up (memory, channel) space is different than truly *new information* being produced. Think of 10k string of random signs vs 10k 0s.
- Transforming what's 'already in there' does not qualify as generative either. Otherwise it's trivial. Every new output would be 'generated'.
- Different cases of *faux*-generation: copying, translation, integration, downsizing.


Generative models

- Generative model—statistical model of the process of how sensory data are generated from a set of hidden causes.
- Functionally different than more classical discriminative models.
- Explanatorily crucial in predictive processing - top-down predictive architectures. (Clark 2015)
- But also semantic pointers! (Thagard & Stewart 2014; Eliasmith 2013; Blouw et al. 2015)


Autoencoders as an example

Autoencoders work by *encoding* unlabeled data into a **compressed representation**, and then *decoding* the data back into its original form. “Plain” autoencoders were used for a variety of purposes, including reconstructing corrupted or blurry images. *Variational* autoencoders added the critical ability to not just reconstruct data, but to output variations on the original data. (Martineau 2023)

(a) Autoencoder




(b) Predictive Coding



Ok but so what?

- We argue that information generation can be seen as production of sensory representations using internal models. This is achieved by *reversing* the process of representation learning, which is projection from sensory inputs to internal models. (Kanai et al 2019:4)
- So it is a decompression from the latent space.
- But it's not a regular decompression as it is merely an expansion of what's already there.
- Think zipfiles then compare to upscaling. Operating solely on the file itself only gets you so far.

Upscaling example



The gist

- However, aren't all these just mathematical functions (albeit fancy) that map input to output?
- How to differentiate biased processes from legitimate inferential generativity? (Orlandi 2014)
 - Real informational differences, not 'explanatory gloss' (Cao 2020)
- Proposal: previously stored information mobilized for generation of a new batch of information (variation on Boghossian's 'taking condition?')
 - How it's implemented - tbd. (weights. connections. other types of memory - e.g. molecular; Colaco 2022)
- This stored information has to correspond to some regularities that enable that filling-in (as in: gist+completion paradigm; fuzzy-trace theory type (Reyna & Brainerd 1995; in AI research: decompression from a seed sampled from more latent, low-dim space)
- **No information without (previous) information, i.e. you need (stored & regimented) information to (internally) generate new information**

The conundrum

From the standpoint of algorithmic theory of information what's 'extendable' (intuitively: unfurled) from the regularities is not really new information.

Any such 'filling-in' according to patterns is less informative than random noise generation (or even not at all?).

But it clearly seems like generativity both intuition-wise and on account provided here. How so?

Two conceptual ways out: 1) brute generation $\neq$ actual generativity; 2) generativity can be assessed for the system (if the relevant information encoding regularity acted upon is actually represented *in* and *for* the system).

Let's combine them!

Semi-conclusions - desiderata towards a definition

- Lost (or not had in the first place) information has to come from somewhere
- Coming back to the idea of inference as filling in the gaps.
- **Takes advantage of regularities** nonetheless (prediction based on the underlying structure) - just **actively represented and used**.
- Exceeds what's in the signal. Information (at least partly) comes from **storage** inside **compressed representations** which are actively mobilized to perform the task (decompress).
- Production performed *in the virtue of their content*?

**Thanks for your
attention**