Frameworks for Causation

Causality in Complex Systems
25/11/09

Nicolas Brodu

Causality

Aristotle revisited in a Complex Systems context [1]:

Efficient causality

- The notion that somethings brings something else

Material

- The cause is the composition

Formal

- The cause is the structure

Functional

- The cause is the role played in a larger system.

Ex: An alarm clock rings. Why ?

Other notions: control [2], predictability… many more

[1] Emmeche et al, 2000. [2] Boschetti/Gray, 2007.

Causal reductionism

Breaking down to simple causes: with which notion of causality?

Efficient reductionism

- The reason why the reason why... need another notion of causality to break the chain.

Material reductionism

- OK, as simple as saying a brain is made of atoms. But not very helpful.

Formal reductionism

- OK? Says that reality can be described in formal equations. Matches physical laws. BUT: Formal systems are known to be incomplete...
- Questions: are there observables / consequential aspects of reality that are not formally reducible? Does it matter?

Functional reductionism

- The function has meaning only by considering the rest of the system. Causal chains are possible though, as above. Semantic vs syntactic [3]. Dictionary as an oriented graph example (with local ordering of edges).

Other

- Control: Suffers from recursion as above.
- Predictability: Gets diluted with each uncertain prediction. Quantifiable influence?

[3] Pattee, 1995.

Emergence = reductionism failure?

And conversely, would something still be emergent if we could find “the reason why” ? ⇒ Depends on the many definitions of emergence.

Still:

- « The whole is greater than the sum of parts »: failure of the formal (sum) reductionism (material still holds).
- self-organization: failure of functional and formal ones
- weak emergence [2]: formal reductionism with a failure of the functional one.

Pattern: material reductionism holds, formal sometimes, but what we are interested in is usually a functional aspect of some higher-level entity.

[2] Bedeau, 2003

Simulations and computers

Computer program as a formal description

- Can possibly simulate any formally reducible system, power of the generic Turing machine.
- Cannot simulate formally irreducible systems.
- Sometimes the simulation is itself not formally reducible to computer instructions. Ex: physical random number generator, asynchronous threads or network links, user interaction, etc.

Formal irreducibility of a given global observable

- We cannot distinguish between "formally reducible but computationally irreducible" (weakly emergent) and "formally irreducible" (incompleteness).
  - Just take formal statements / computer programs written as strings, one by one in order, and test. Logical undecidability of when to stop: some larger simulation might produce the global observable we wished to reduce.
- Does it matter? Are simulations useless?
  $\Rightarrow$ No, we can still get an understanding of what’s going on.

Understanding a system

We'd like a reasonably concise description of the system, allowing for reasonably precise predictions that can be tested, validated, refuted. Here we start from observed global effects and would like to understand them (reduction approach).

- Ex 1: A straight line pattern appears on screen when plotting the zeroes of the Riemann Zeta function. Not currently a formally reducible statement.
- Ex 2: A given simulation shows points (agents, whatever) appearing in definite higher-level patterns, with noise. Ex: an exponential or parabolic curve.
  ⇒ Make a higher-level law from the patterns and then try to refine with further experiments.
  ⇒ Consider the system as its own little world and apply the scientific methodology.
- Ex 3: Entities with spatial boundaries (groups of agents, patterns). This time the generalization of the higher-level law has to be made both in time (across simulation runs, as before) and space.
- Ex 4: Same as all the above, but with non-deterministic simulations. The observations remain when using physical RNG and distributed computing lacking proper event timing.

Understanding is then better achieved by considering the new level as itself, and applying the scientific method to get a new formal system from scratch at that level, independently of the lower-level formalism [5] ⇒ formal reductionism is useless in these cases.

Not incompatible with trying to build the global function from scratch, ex: MIT "proto" language. It's just that there is no general guarantee the global observations are reducible in the first place! Backup plan: Dealing with good enough approximations that are satisfying for all practical purposes.

[5] Fromm, 2006

Causality as a tool

From the previous part of this presentation:

- We'd like to find the cause for some entity / pattern / relation.
- Functionally defined entities are the hardest to grasp, formal ones might be difficult to deal with, material ones are usually fine.
- Better to consider the system in its own, and apply the scientific method from scratch.
  ⇒ We need instruments, tools, to build the theories specific to a given system. Then try to validate / refute, refine, test predictively... business as usual, what Kuhn calls "normal science" [6].
- Formal reductionism does not matter so much for understanding a system.
  ⇒ Understanding brings its own formal system, related or not to lower level.
- Computers are good tools, even in the face of formal irreducibility.
  ⇒ Algorithms are our generic instruments for investigating complex systems.

Next part of this presentation:

- Tools based on causality, in a computer / formal system framework.
- Using these tools may help us find relations between interacting entities in a given complex system. And then, we have hope for understanding.

[6] Kuhn, 1962

Granger Causality (short intro*)

Two time series, of random variables X and Y, and possibly a vector of supposed explanations in Z.

General Idea:

- Try to predict future(X) from past(X,Z).
- Now repeat with future(X) from past(X,Y,Z)
- If we obtain better forecasts, then Y "tells something" on the future of X.

Shortcomings:

- What if we also get better forecasts reversing the roles of X and Y?
  ⇒ Not necessarily wrong, we just have not captured all dependencies and feedback effects. Perhaps there are unseen common factors as well.
- Original G-causality is based on linear regression. But the idea may be extended: non-linear predictors, other measures of information gain and prediction success, multivariate cases, etc.
  ⇒ The "causality" we get depends on how we compute it. This is not surprising, and it's anyway not worse than other notions of causality.

Usage as a tool:

- Granger causality may help building causal chains / graphs between entities in complex systems. It is an aid to the experimentalist in building the understanding of the system.

* We'll have more complete descriptions in this workshop, there's no use spoiling them now!

Causal and Bayesian networks

Graphical representation of causes and effects

Joint effect

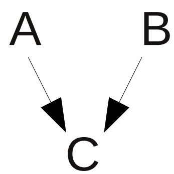

Common cause

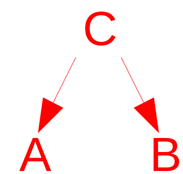

Chain

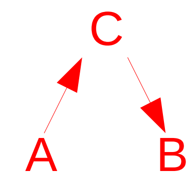

Arrows here represent a causal influence, whatever that means given the definition chosen for a causal influence.

Graphical representation of a joint probability distribution

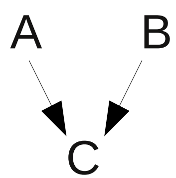

$P(A,B,C)=P(A)P(B)P(C|A,B)$

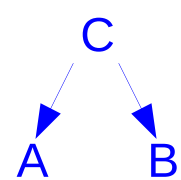

$P(A,B,C)=P(C)P(A|C)P(B|C)$

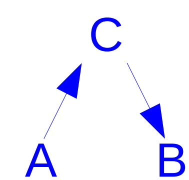

$P(A,B,C)=P(A)P(C|A)P(B|C)$

Arrows represent the conditional independence structure of the data.

Can we identify both?

Not in general: the BLUE graphs are equivalent, the RED ones are not.

BUT

There are techniques to help distinguish the blue graphs. Ex: Intervening on C would have an effect on A and B or just on B depending on the causal graph.

Causal and Bayesian networks

Idea: causes presumably do not depend on their effects:

- A true causal structure imposes constraints on the possible decompositions.
- If these constraints are observed in real data, we may hope to reconstruct some of the causal graph. ⇒ A limited but useful definition of causality.

In practice, given that:

- Markov Condition (MC): Parents shelter variables from their ancestors.
- Minimality: no subgraph has the Markov Condition
- Causal MC: The graph matches its causal interpretation.
- Faithfulness: All arrows are necessary for Causal MC to hold

Then:

- Equivalent forms were found and classified [6], where equivalence means all graphs in this class match the data equally well.
- Algorithms exist to infer these forms from the data, given the above conditions.

However:

- How to assert the conditions hold? Ex: hidden variables.
- What if the variables are not on the same level, “emergent” in some fashion ?
- What about time, and cause preceding effects?

In any case, this tool is powerful, well-established, and renders some form of causality

[6]

Why Probabilities?

Why try to make these networks (causal, Bayesian) match in the first place?

- Explanation 1: “Correlation is not Causation” is a famous warning in statistics, rising the immediate question: “What then, is Causation?”
  ⇒ Generations of traumatized statisticians have and will continue to investigate the subject, more so than in other disciplines, which explains the use of statistics!
- Explanation 2: Probabilities interpretation as relative frequencies of occurrence, which we presume are different for an effect when one of its cause is involved.
  ⇒ Are there definitions of causality without any effect on frequencies of occurrence? Are they interesting?
- Explanation 3: Deterministic causation is too strong for many practical use. Ex: “A ⇒ B” as a logical implication statement means that whenever A is true B is also true. But for problems like health care, like “smoking causes cancer”, we do not mean that cancer must follow absolutely the act of smoking, only that the probabilities have changed.
- Complex systems specifics (functional definitions of entities, building formal systems, large number of elements...) push to statistics as well.

Other formal tools:

Kolmogorov-Chaitin complexity, structure of graphs (small world networks), indicators of order and disorder and the edge of chaos, attractors and measures of stability like Lyapunov exponents and the like, etc.

Causal states and ε-machine

Predictions in time series

- A={a,b,c} alphabet. A* strings.
- X=...abbbacacabcaac $= s_{t\leq p}$. $Z = s_{t>p}$ past and future strings
- Finite-memory Markov condition: if we consider than an horizon h in the past captures all predictive information: $x = s_{p-h < t\leq p}$. Idem for future.

In a physical context (with information transfer speed)

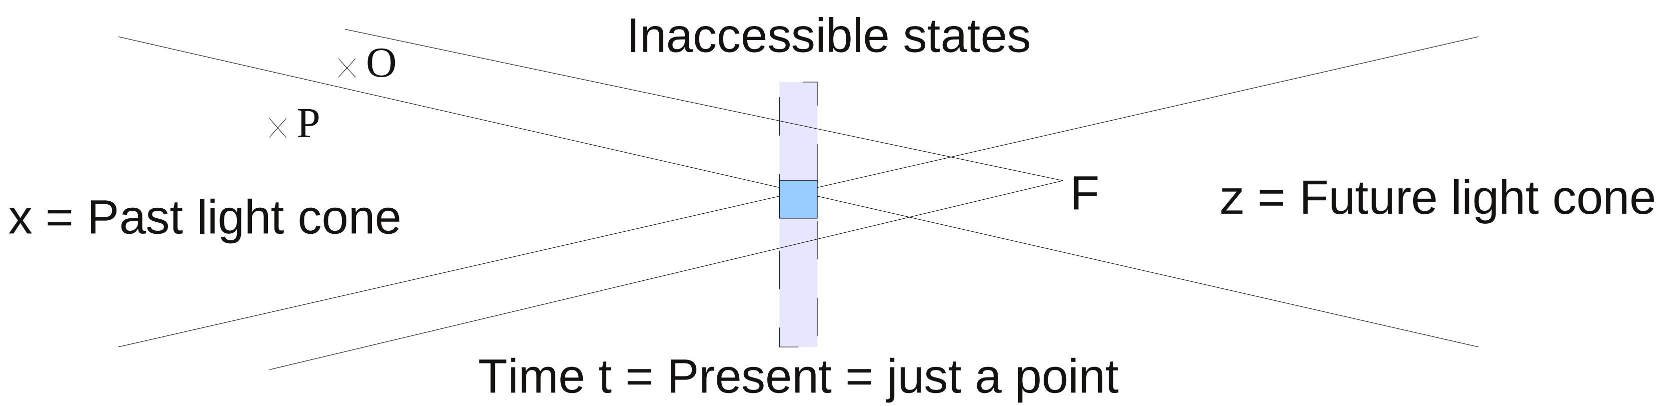

Light cone definition: all system states in time that may have a causal influence on the present state, for x, or that can be causally influenced by the present state, for z.

More generally: Consider a space X. Z a space of predictions from X.
We try to capture p(Z|X) in order to characterize what may happen.

Causal states and ε-machine

Refs: Shalizi [7], Crutchfield, Young, and more [7 for pointers].

Idea: cluster together all configurations that have the same distribution of futures.

Formally: $\sigma(x) = \{ w: p(Z|w) = p(Z|x) \}$ is called a Causal State

Note: In the case of temporal systems, we cannot distinguish x and w in the same state by new observations ⇒ equivalent pasts.

With symbolic series, $x_{t+1} = a x_t$ with $a \in A$ a symbol.

But $\sigma(x_{t+1})$ exits too ⇒ The σ form a deterministic automaton, called the ε-machine:

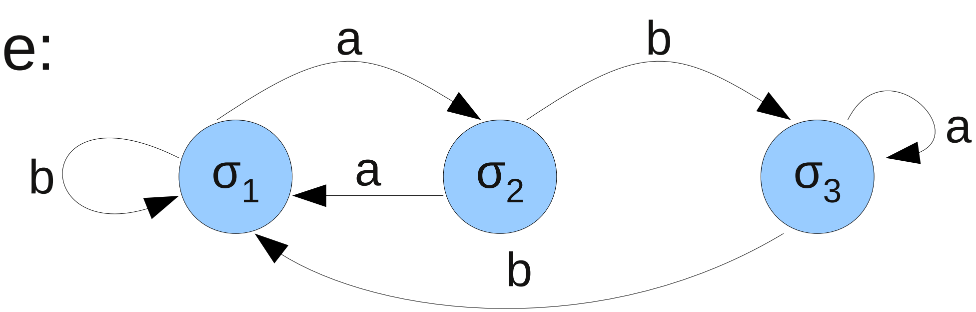

The ε-machine is a Markov model of higher genericity than finite-memory Markov Chains (they can deal with sofic shifts)

 [7] Shalizi, 2001

Causal states and ε-machine

Causal states and causality:

- All we have is a misleading predictive view on the system:
  - Ex: Coughing and Sneezing are good predictors for a cold. But they are the effects, not the cause!
- But when restricting to past / future predictions:
  - Then presumably the causes are good predictors.
  - Are there better predictors than the causes?

Causal states for complex systems:

- These are the optimal and minimal predictors within the class of deterministic automata ⇒ we get a quantifiable measure of how complex / difficult it is to make predictions in a system. $C = MI(x, \sigma(x))$ = entropy of the causal states in the discrete case.
- A measure that is low both at the totally ordered and totally random ends, while getting high only in between.
- Computable from data ⇒ a good tool to check where there are interesting things happening in the system !

An extension of my own (ongoing work):

- Give a utility/cost $U(y,z)$ to predictions y when z actually happens.
  ⇒ A priori knowledge can be introduced in the system on top of the ε-machine.
- Leads to decisional states, and a measure of how difficult it is to take a decision.

Conclusion

Causality in complex systems:

- Both a core issue: The reason why some function appears, “emerges”, at a high level.
- And a practical tool: Causal inference and model building techniques help identify entities of interest and their interactions.

Philosophically:

- Too much to say! This short though-piece is just a small part, dealing especially with complex systems.
- Introduction of the “formally irreducible” kind of emergence. Why it does not matter is the basis for a system exploration methodology (whether it will prove fruitful or not is another issue...)

Practically:

- Formalisms exist where we can measure/capture at least some kind of causality. 3 ex: Granger Causality, graphical models, and ε-machines.
- More tools can be used, the goal is to understand and formalize what's going on in a given system.

Thoughs to share:

- Probabilist causality is good, but in C.S. we deal with more than just well-defined random variables. How to extend the above frameworks (or build others) so as to deal with a degree of uncertainty on the identity itself of the entities we manipulate? How can we deal with collectively defined entities, possibly only valid in the limit to large numbers, and how to apply the formal tools to them? Are there better notions of causality?
- How to automate the process (formal epistemology) in order to build a more or less ready-to-apply toolbox for the investigation of complex systems?

Thank you, questions / comments are welcome!

References

[1] Claus Emmeche, Simo Køppe, Frederik Stjernfelt (2000) Levels, Emergence, and Three Versions of Downward Causation. In Downward Causation, Minds, Bodies and Matter, p13-34, Århus: Aarhus University Press.

[2] Fabio Boschetti, Randall Gray (2007), A Turing test for Emergence. In M. Prokopenko (ed.), Advances in Applied Self-organizing Systems, Springer-Verlag, London, UK, 2007.

[3] Howard Pattee (1995) Evolving self-reference: matter, symbols, and semantic closure. Communication and Cognition - Artificial Intelligence, 12 (1-2), 9-28.

[4] Mark Bedau (2003), Downward causation and autonomy in weak emergence. Principia 6: 5-50.

[5] Jochen Fromm (2006), On Engineering and Emergence. Unpublished preprint, arXiv.org:nlin.AO/0601002

[6] Thomas S. Kuhn (1962), The Structure of Scientific Revolutions. University of Chicago Press. Reprinted in Champs, Flammarion.

[7] Peter Spirtes, Clark Glymour, and Richard Scheines. (2000) Causation, Prediction and Search, Second edition. Cambridge, MA: M.I.T. Press.

[8] Cosma Rohilla Shalizi (2001), Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata. PhD dissertation.