Complementing Language Models with Persistence Diagrams

Raouf Dridi, Steve Reinhardt
steve@xfr.ai
Transform Computing, Inc.

XFR
Transform
Computing

We have initial evidence that combining persistence diagrams of energy grids with a language model delivers analytic insights beyond what the language model provides.

XFR
Transform
Computing

Context: Language Models are Untrustworthy

- Generative AI as implemented by current language models (LMs) is prone to errors (“hallucinations”)
- GAI is not able to sanity check via logic or math/physics knowledge, hence viewed as untrustworthy
- Several of the industrial use cases predicted to benefit most from GAI need high confidence in the results, e.g.
  - **Energy**: confident control of operational systems
  - Healthcare: computer-aided reading of medical images, computer-aided diagnosis
- XFR sees opportunity for explainable AI (XAI) addressing these shortcomings
- For the foreseeable future, any use of XAI for real-world use cases will be “computer-aided <task>”, not complete automation

XFR
Transform Computing

Context: Analysts Could Often Benefit from Advanced Math They Don’t Grok

- E.g., topological data analysis’ persistent homology (PH) capability is effective at seeing past noise, but Betti numbers and persistence diagrams are not easy to map to a given domain.
- LMs are good at translating. If they can translate from a given domain to a particular math topic and back, that could deliver the math value to users who don’t understand the math. This translation could be high value.

Thesis: Persistent Homology Can Provide A Concise Sanity Check on LM

- (We are not pursuing building a new type of learning mechanism from the bottom up.)
- LMs can only predict in a language sense, not in a math/science sense
- Persistence diagrams (PDs) concisely describe the “shape” of data in a noise-resistant way
  - And can be represented as text for digestion by an LM
- Can LMs+PDs deliver insight into numerical dynamics, and hence a degree of explainability?

Target: Renewable-energy Mesogrids

- Shift from fossil fuels to renewable energy is an urgent societal imperative
- Traditional grid is static in key dimensions where Smart Grid will be dynamic
  - Unidirectional energy flow v. bidirectional
  - Renewable generators are often intermittent (solar, wind, tidal)
  - Rotating mass to counteract frequency/current drift *v.* need to actively respond
- Number of grid nodes increasing radically – each solar panel, windmill, and battery; strong need to identify/respond to emergent issues
  - Increasing complexity is overwhelming conventional analytics
  - Need advanced and innovative capabilities
- Regulatory and commercial environment will be dynamic well into the future, so adaptability essential

XFR
Transform Computing

Target: Renewable-energy Mesogrids (cont.)

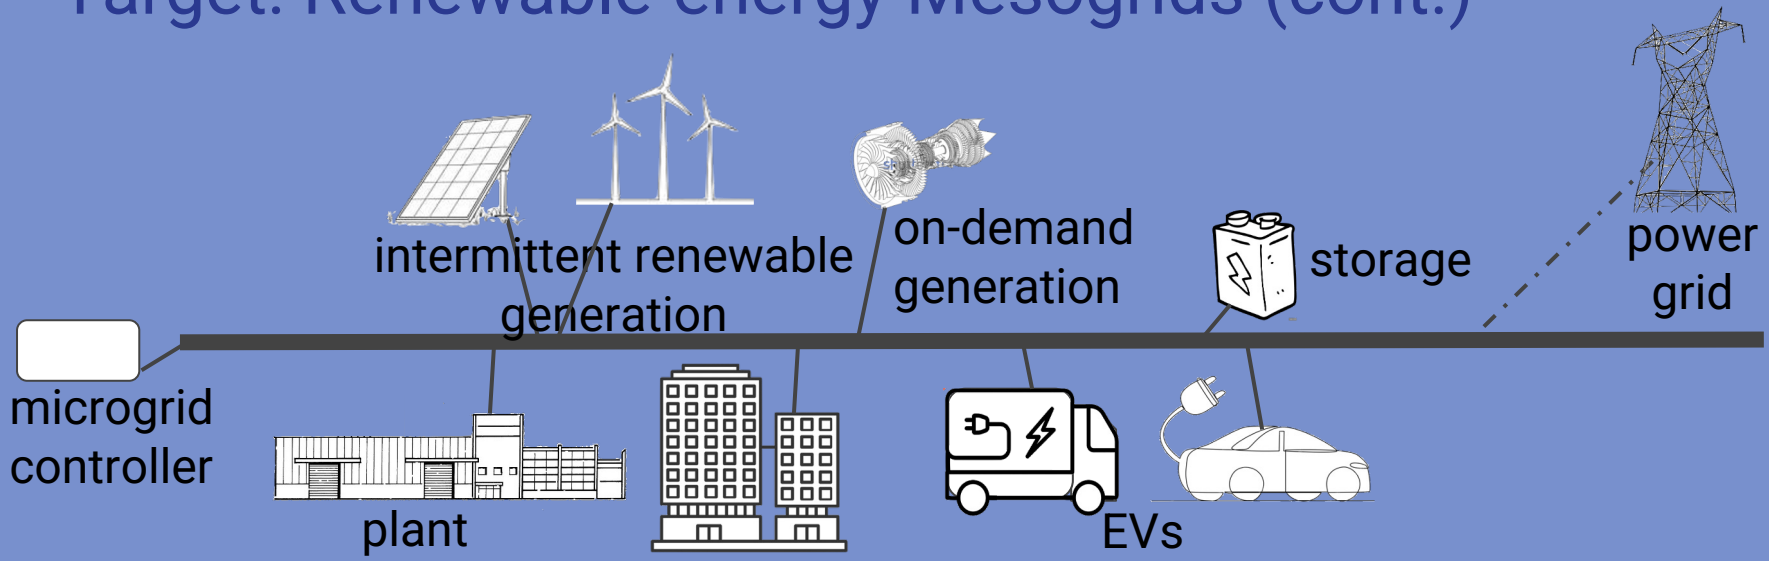

- XFR targeting **mesogrids** == subsets of the grid controlled by one org (e.g., producer, utility, industrial consumer); architecturally are **microgrids**
- Microgrids are key bridge enabling semi-autonomy for faster grid evolution
  - Ability to disconnect (“islanding”) and reconnect to the main grid is important

XFR
Transform Computing

Responding to Emergent Mesogrid Issues via LM+PDs

- Today, microgrids are limited by size, richness of config, and effectiveness, largely due to inadequate analytics
- Difficult to detect key microgrid issues – e.g., real-time power balancing, adaptive protection settings, and islanding transitions

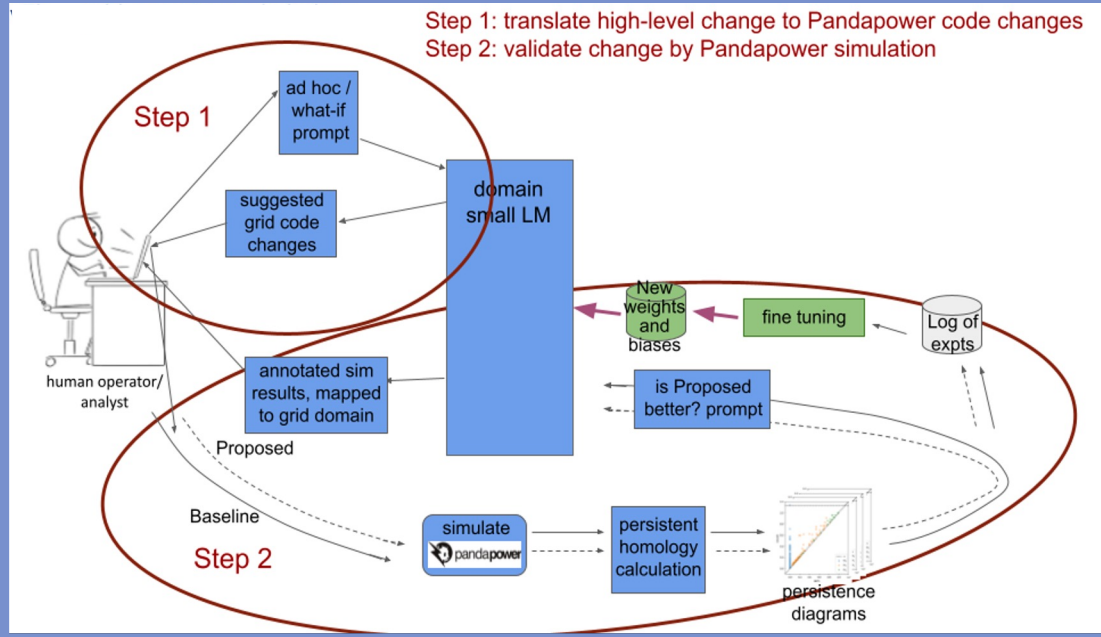

- Idea: use persistent homology (PH) to detect emergent issues and use LM to map PH results from math-speak (“Betti numbers”) to grid-speak

XFR Transform Computing

Return to Overview slide

Implementation Details….

Caveat: to date, zero attention to user interface

Lessons

- LMs (e.g., Claude) are reasonably effective at mapping an electrical attribute to its energy-grid meaning
  - Though appear to be more adjacency focused than topology focused
  - Loading percent and voltage appear to work well; susceptance (== 1/reactance) must be inverted to be a distance, which creates numerical issues
- LMs effectively map persistence diagrams to their energy-grid meaning
  - At least for first-level issues, such as connectivity and redundancy
  - Or cybersecurity meanings, as shown by DARPA IMPAQT project
- LMs are surprisingly effective at generating code from natural-language guidance to achieve a particular goal
  - “I want to increase the grid’s resilience in the vicinity of X”
  - LM learned to generate code changes as an ed script 😄

XFR
Transform
Computing

Practical Software Engineering Points

- [ephemeral?] Prompt-/response-/context-window-size limits were issues but appear to be going away
  - "If you need to work with extremely large matrices (e.g., thousands of rows and columns), ..."
- Used LangChain so ~LM-independent
- Generate code separately and tell LM to import code from a given file name or generate output to a given file name
- “Follow-up” capability has been useful, at least during development stage, since LM often doesn’t get request exactly right

Technical Strategy

- LMs are great as far as they go
- Topological data analysis (TDA) and its persistent homology (PH) capability are effective at surfacing *persistent* characteristics of complex data, which enables humans to see past noise in the data. Good complement to LMs. Quantum-capable, so will benefit from quantum computing's radical reduction in power consumption.
- Logic theory and systems deliver rigor and insight, complementing LMs well. Precise natural language <-> logic language mapping important.
- **Explainable AI (XAI) is the top-level concept that unifies these, though not all use cases require/deliver XAI**

XFR
Transform
Computing

Next Steps

- Workflow works end-to-end
- Needs further calibration by mesogrid experts
- Scale to real-world-useful mesogrid sizes
- Are mesogrid attributes other than load percent also valuable?

XFR
Transform
Computing

Summary

- LM enables a domain expert to exploit the power of persistent homology to operate energy mesogrids (without knowing math details)
- Grid calibration/tuning remains
- Energy grids have computer networks inside them; same approach appears useful for cybersecurity

- This approach can be used for other math techniques (e.g., symbolic logic)
- This approach can be used for other domains besides energy grids

XFR
Transform
Computing

Backup