Engineering Ecosystems with AI by Alex "Sandy" Pentland, 9/15/23

Matthew Webb

Part I. Literature (in order of appearance during the presentation)

Yuan, Jahani, E., Zhao, S., Ahn, Y.-Y., & Pentland, A. S. (2023). Implications of COVID-19 vaccination heterogeneity in mobility networks. Communications Physics, 6(1), 206–209. https://doi.org/10.1038/s42005-023-01325-7

Krumme, Llorente, A., Cebrian, M., Pentland, A. S., & Moro, E. (2013). The predictability of consumer visitation patterns. Scientific Reports, 3(1), 1645–1645. https://doi.org/10.1038/srep01645

Chong, Bahrami, M., Chen, H., Balcisoy, S., Bozkaya, B., & Pentland, A. "Sandy." (2020). Economic outcomes predicted by diversity in cities. EPJ Data Science, 9(1), 1–19. https://doi.org/10.1140/epjds/s13688-020-00234-x

Yabe, Bueno, B. G. B., Dong, X., Pentland, A., & Moro, E. (2023). Behavioral changes during the COVID-19 pandemic decreased income diversity of urban encounters. Nature Communications, 14(1), 2310–2310. https://doi.org/10.1038/s41467-023-37913-y

Caros, N. (2023). Leveraging spatial relationships and visualization to improve public transit performance analysis. [Doctoral thesis, Massachusetts Institute of Technology]. DSpace@MIT. https://hdl.handle.net/1721.1/139338

Krafft, P. M., Shmueli, E., Griffiths, T., Tenenbaum, J., Pentland, A. (2021). Bayesian Collective Learning Emerges from Heuristic Social Learning. Cognition, 212. https://doi.org/10.1016/j.cognition.2020.104469

Adjodah, Leng, Y., Chong, S. K., Krafft, P. M., Moro, E., & Pentland, A. (2021). Accuracy-Risk Trade-Off Due to Social Learning in Crowd-Sourced Financial Predictions. Entropy (Basel, Switzerland), 23(7), 801–. https://doi.org/10.3390/e23070801

Dubey, A. (2020, November). Kernel methods for cooperative multi-agent contextual bandits. In International Conference on Machine Learning (pp. 2740-2750). PMLR. https://doi.org/10.48550/arXiv.2008.06220

Girotra, K., Meincke, L., Terwiesch, C., & Ulrich, K. T. (2023). Ideas are dimes a dozen: Large language models for idea generation in innovation. https://dx.doi.org/10.2139/ssrn.4526071

Part II. Recent News

Ratti, C. & Picon, A. (2023, September 14). *AI is coming to our neighborhoods and will show us the future of cities*. The Boston Globe.

https://www.bostonglobe.com/2023/09/14/opinion/artificial-intelligence-cities-predictions/

MIT scholars awarded seed grants to probe the social implications of generative AI. (2023, September 18). MIT News. https://news.mit.edu/2023/mit-scholars-awarded-seed-grants-generative-ai-0918

Lomas, N. (2023). Poland opens privacy probe of ChatGPT following GDPR complaint. TechCrunch. https://techcrunch.com/2023/09/21/poland-chatgpt-gdpr-complaint-probe/

Part III. Questions

Jinhua: Many people ask the question about will AI play a convergence role versus divergent role in terms of a different group of people. You made the comment that LLMs may help the medium skilled more than the high skilled workers. I want to broaden this question to say what are your thoughts about AI in general on the impact on society in terms of convergence versus divergence. A second, related question is the fact that LLMs helped medium skilled workers more than high skilled workers, but it is after the fact that we empirically tested and determined the result. Is there a way we can engineering a system like LLM with that purpose in mind? Can we design an AI system to achieve that conversion role beforehand? Is it possible or are we just at the mercy of the whatever comes out of this?

Pentland: What I see is that there is an enormous race not to make bigger models, but to build specialized, quite smaller models. So, in their programming, what is the data they're going to use to train on? It's going to be from experts, and so the things that come out of that model are not going to help experts all that much, but they are going to help, according to the evidence, mid-skilled programmers. And I also showed about marketers, with that report out of Wharton. People will use this technology in bad ways, but people use all sorts of technology in bad ways. It's the people not the technology. So, what we want to do is try and figure out how we can use this technology to help people. What I'm suggesting is given what we know about how people make decisions and how they learn that helping them by presenting them with the common sense what people normally do is a great strategy. It leaves the person in charge so a lot of the ethical questions are minimized and it seems to be a major way where we can result in upscaling of people and perhaps transition them people from mid skill to high skilled jobs, which is what we need as AI deploys. We need people who are able to deal with this stuff. I would suggest that what we have to do is focus there. Is how do we build things that help mid skill people to go to high skill by telling them what "everybody knows." In companies this is called the company culture. Like if you work for Ernst and Young, you know what are the ways to approach things. But if you're a new guy at Ernst and Young, you don't know that and this could be very useful for that.

Jinhua: You mentioned continuing education, which we've been calling for for multiple decades saying we should have lifelong learning and that school is just the starting point. But based on your empirical findings, we largely failed to achieve any result. Do you see this, LLM or AI can be an opportunity to really change that failure? That's actually related to the single project we've been collaborating on for how AI can improve human capital. Can you give some thoughts on how AI can play a better role in continuing education?

Pentland: Well, this is embedded continuing education so if you're a new person in a business or a medium skilled person trying to become more skilled this is something you can use to tell you what the common sense is and what other people do. You will learn as a human over time and you get the skills of how to do the job by looking at the things that are suggested that other people do and considering them for your situation. Then you can reflect on these too. You can say, “well, I think this is the best choice” and you can ask the LLM what other people will say. And it'll give you a sort of, cultural average in terms of a response, which you can choose to ignore, of course, right? It helps you guide your thinking and your learning and that idea of having embedded learning has pretty clearly got some advantages over having to take time off work to take separate classes.

Jinhua: You mentioned in the first half of the presentation that your lab has done a lot of work in in terms of this social interaction in the physical sense. For instance, people living in one place, travelling in another neighborhood, and maybe shopping another neighborhood. I saw the fact that COVID really hurt that quite a bit. One way to boost the social interaction is to try through physical infrastructure. Then in the second half you presented AI or a large language models as another alternative way to boost the social exploration. So how do you see the two potentially interacting with each other or helping each other or substituting each other between the physical interactive way versus the AI based way of hosting social interactions?

Pentland: Well, I think the human interaction is the primary thing because there's all this stuff that isn't in language: body language, attitudes, emotions. Some of those have echoes in language, but they're not there completely. We have a real need for human interaction, so when we segregate our communities you can see all the bad things that happen. One of the main things that happens is distrust. And distrust results in polarization and we all can see a lot of that. At the moment, it's probably one of the major problems that we face as a society is distrust between different strata, different communities, etc. So, the physical is absolutely critical. On the other hand, you also have this more sort of cognitive skills and experience that you've had and if you can accelerate that, that strikes me as a good thing. Like if you go to the party here's the things you ought to remember that they do at that sort of party. Like, don't wear that sort of dress or don't say these sorts of things - the social norms. So, if we can build things that do that, I think that will help the situation. I feel looking at the, comments here that you can do a lot of bad things too. Let's not do that, okay? That's the correct answer to all of those comments. Don't do that. Instead let's focus on the good things and then we will figure out ways to, discourage, prohibit, etc, the bad things. I do a lot on that. I was just that the UN board meeting about how to do things like that. One of the primary things is that you have to keep track of what this stuff is doing. We don't do that in our society. We don't know if the algorithms are plus or minus. We don't know who they hurt and who they help because we don't keep track. We have to keep track and that has to be visible to other people so that we can find the bad things and we can figure out how to fix them. If you don't know what's going on, you're never going to fix it and that's another sort of major source of problem for society; a lack of transparency and accountability.

Jinhua: Thank you. The last question from me, and then I’ll pass it to Bhuvan for the curated audience questions. Last night you recommended the paper Optimized Human-AI Decision Making for me to do as a pre reading. In that paper you mentioned exploration versus exploitation, particularly, the difference between individual decision making versus the group decision making. It seems that AI can help the individual in a very different way then how AI can help with a group. For example, you mentioned that as a group, when you set off, choose the

maximum likelihood action, you could do distribute your action into the frequency of action proportional to the likelihood so that you have enough exploitation but also you have enough room for exploration. My first question is how do you see the difference between AI helping an individual versus helping a group? The second question is that in the paper you mentioned that the human-AI system design has an optimality property where the demand is expressed as a learning outcome. For example, you mentioned finance. So, what about the demands that do not have such explicit learning outcomes? How do we think about the better design of this human AI system?

Pentland: The fundamental thing here is that most of our culture focuses on individuals making decisions essentially by themselves as rational individuals but that rarely describes humans. You can go to school for 20 years and learn to do that. It's hard. That's why it takes 20 years. In fact, most of our decisions are culturally bound, and that's a good thing because that means that we can learn from the experience of others. In terms of sampling theory, optimal estimation, or any of those sorts of classic ways of looking at determining best strategies, learning from others is a huge win. It multiplies your abilities by orders of magnitude. And that's the sort of social thing. In any particular community, the learnings of the community, we often give it names like culture. This is the way we do things here. and it can be wrong, but it is sort of the learned answers of this group of people. And, you know, approaching that as opposed to changing the individual is often something that's easier because you're adding other opportunities to the community. People have that rationale, and they will tend to make those choices. In terms of objective outcomes where you can see things quickly in a very hard way, learning stuff happens pretty well there because you can see the result. One of the major problems we have as a society is that we do things that look good in the short term, but in the long term they're terrible. Climate change and Finance are good examples of this. There are many things like this. Typically, the only way from a mathematical point of view to do that is to aggregate over larger and larger numbers or more and more experience. So, if it's just me making decisions, I might do stuff that looks really good, but then it kills me later. If I have samples from hundreds of thousands of people, I can say for some of these people, it looks pretty bad and the trend is in the wrong way. So, that's like learning from a broader community, which we, 10 not to do because we don't have transparency and accountability. We don't know outcomes of actions in some way that is reliable and truthy.

Bhuvan: Thank you, Sandy, for that fascinating talk. The chat as always has exploded with comments and questions and I'll be sure to send you a copy so that you can go through it before your talk later on. I'm going to combine two threads of questions into one. You mentioned how most LLMs have been trained on Reddit and that's probably not the best source.

Pentland: Yeah, it's a joke also, but true.

Bhuvan: There is this whole idea of bias in AI, and particularly with large language models, on what data set have you trained? You gave the example that company cultures can be transferred in terms of workforce training and development. Company cultures have a legacy of racism and gender bias built into them, so will AI perpetuate these cultures? Is this really a social benefit?

Pentland: First, the current generation of LLMs are trained omnivorously on everything including Reddit and that's probably a major source. We have a project called the Provenance Project, which is going through and surveying open data sets about things that you care about like is this truly used by third parties, or was this done under human subjects' approval? This gets much cleaner training than just having to go over Reddit. I think that's the trend that will

happen. Everybody is racing to establish good training corpi that can be used in a legal way and can be used in a way that is authoritative. It won't cure all the problems, but it'll cure a fair amount. And then the company culture and racism, so forth. That's auditing, transparency, and accountability. The real problem is that we don't know until somebody collects the evidence, which is very difficult, and makes the argument, which is very difficult. There ought to be standard ways of monitoring these things and computers could do that. It's not hard. It's not expensive, but then there would be a public repository where you could say, okay, so here's the company policies. Is it racist? Oh, statistical test. It should take you a good solid 20 seconds to answer that action as opposed to 5 years of your life. I remember a talk I was at in Oxford where a leading person in this field got up and talked about bias and everything and the justice minister of Kenya said, "what you say may be true, but have you seen our current system?" The humans in our system (like judges) are incredibly biased and we don't call them to account either. I'm a big advocate of we ought to audit everything. We ought to be able to get feedback about actions versus sensitive categories and about outcomes both short and long. Then we'd have a hope of being able to find policies and actions that are good in both the short and the long term. Until we have that sort of auditing, that sort of data I think it's all heuristic. I know that's not terribly popular, but that's why the sustainable development goals at the UN have data metrics. It is because unless you have data metrics that people agree on it's all a bunch of hot air.

Bhuvan: Initially in your talk you mentioned that LLMs may be helpful to medium skill workers. Here there are two aspects. One is helping them with discrete tasks versus insights. So, how can AI enhance the medium skilled workers with insights, or can they not in enhance them with insights and just help them, but not take their jobs in the future? So what is the balance there? Enhancing their jobs now, but possibly take over their jobs in the future or can AI actually give them the insights they require to move up from a medium to a high skilled job?

Pentland: I don't have all the answers, but what I believe is that we can design things that will help Lower skilled workers to become higher skilled. That's a major step towards solving those problems. Will it solve it all? Probably not, but if we have less segregation in our society, so there's better spread of opportunities, and we have a continuous upscaling going on, it seems like we're likely to be better off than if we don't do those things and we continue in sort of a segregated, siloed society. I see there's a lot of things in the chat about privacy and yes you can do stupid things. Don't. One of the major threads of what we do is around privacy and how you do auditing. Where you don't give up personal data at all. I was the one that led the discussion at Davos that turned into GDPR. Privacy is sort of at the core of what we do. Take a look at the sort of things that are happening there. It's now possible to audit things in ways that were impossible only a few years ago in terms of privacy and in terms of security. There are some inevitable trade-offs. If you're going to hold people accountable, then you have to know who to be accountable, but one can make a judicial path for that. Without data, without knowing what's happening, it's going to be very hard to do anything sensible.

Bhuvan: It's great that you brought up the point of privacy as well. There was one comment that mentioned that they'd be curious as to what percentage of conversations that happen across the globe on any given day are digital. There's this whole aspect of your car listening to you all the time, Alexa and Google Home listening to all these conversations and it being transferred. What are your thoughts on privacy in the GDPR in EU? Are these really being used? What are the ways that one can have a hard turn off button on these devices?

Pentland: Well, currently they're not being used? First, some of these things are relatively new from a technological point of view so that's not too surprising. Some are the economic incentives of the big companies are not aligned with our incentives. One of the things that we try to work on is distributed systems where you control your data. You can work with your community to get insights that you want and you can choose the community. If you look at transformers.mit.edu. You can see a lot of the sort of stuff that we're thinking about in that sort of area because people want to learn from their friends, but they don't want to post it on Reddit. There is this delicate dance you have to do with data and privacy and accountability. Good news is that the mathematical tools are being, and have been largely, developed to do a lot of this. There still is work to do. The bad news is that the current systems don't use these and we have to cook up ways to "encourage" those things, which is going to be a combination of law and economic incentives of some sort. You could have a long conversation about how you change the economic incentives, but people are thinking about it. They're doing it. We're involved in that conversation. It's not being ignored; it's just a difficult conversation. Let me just put something out there so people understand. General Secretary Xi of China, the largest voice for Marxism, recently said, "data is a new primary means of production along with capital and labor." If you think about that, what he's saying is that classic Marxism is done. It's now not a battle between capital and labor. It's a battle between data, capital, and labor. That gives you a sense of the magnitude of this problem. If you look at what society did with capital and labor, it took a century or more, for instance, to form labor unions, to pressure companies, to establish, principles, and to get laws enacted. The same thing with capital with agricultural banks and credit unions. It's not a fixed thing. It's not like you can do it once and it's done. It evolves over time. Currently we're in a new evolutionary phase of labor and a new evolutionary phase of capital. The problem with data is we don't have any institutions. We don't have any norms for it, basically, it's new. We're back in the robber baron era of capital. We're back in the early industrial age where kids were working 14 hours a day. That's where we are with data? What we have to do is develop the right sort of institutions to be able to deal with this now. A critical element of society, something that you take seriously. It's not just an extra thing or an irritant. It's a major part of society.

Jinhua: Related to that, yesterday the Senate Majority Leader Chuck Schumer tried to organize a group of Tech Companies to think about AI regulation and the role the federal government? What is your view on the federal government's role or general area should they be in? More specifically, related to the fact that AI can be used as a tool to boost the social exploration, is there anything that's necessary from the policy perspective to boost that? Any concerns that we need to address?

Pentland: We work with the EU on their AI regulation, and I work with Brookings, which has helped shape the Biden Administration. I work with other countries too. There are lots of details in there, but the primary thing that I come away with is we don't really understand all the risks and dangers. We don't really understand what are the good things? It needs some more exploration, but that has to be safe exploration. That's why I'm such a big fan of auditing. Perhaps the audits are not public, but it ought to be the case that if you worry about something you can get a court order and an examination of what it's doing against any sort of particular problem. That should be a matter of minutes, not years; of one dollar or two, not millions of dollars. In our legal system a lot of big problems are noticing that there's something wrong and then doing discovery, which is a legal term, which is hugely expensive and slow, and then

litigation. We could take that and turn that into something that you know your average person could do for a $1.95 in half an hour. That would be progress. Fortunately, that's part of the vision that people have. They don't do quite as far as I do. But I notice that the first actual operational AI regulation out of Singapore, is that they are auditing the AI. If you want to release a product and you claim it does X, you have to prove it before you get licensed to do it. Sounds like not a bad idea and the only thing that's wrong with their regulation is that I want them to audit it every month because the world changes and, things go off the rails.

Part IV. Summary of Memos.

Themes

1. Appreciation for the potential use of AI tools to democratize information when used responsibly.
2. Concerns about Pentland's assertions that the "upskilling" of "mid-skill" workers to "higher skill" professions is inherently beneficial or desirable. Audience members asserted all work should have dignity and it is inherently difficult to categorize and compare vastly different professions.
3. Concerns about the use of AI tools and their effects on society. Specifically, concerns about possible bias inherent in the tool or training dataset and how this could have an opposing effect to the social exploration benefits that Pentland champions.
4. Inspiration from Pentland's example of building transportation networks to promote social mixing, exploration, and community interaction.
5. Agreement on the need to cautiously move forward with an emphasis on auditing and transparency in AI tools.

Part V. Other Information

[Paper describing the use of liquid neural networks mentioned by Professor Pentland] Chahine, M., Hasani, R., Kao, P., Ray, A., Shubert, R., Lechner, M., Amini, A. & Rus, D. (2023). Robust flight navigation out of distribution with liquid neural networks. *Science Robotics*, *8*(77), eadc8892–eadc8892. https://doi.org/10.1126/scirobotics.adc8892

[Website for the startup using AI to manipulate cattle microbiome and reduce methane emissions] https://www.metha.ai/

[Paper Jinhua references in his final question] Pentland, A. (2021) Optimized Human-AI Decision Making: A Personal Perspective. In Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI '21). Association for Computing Machinery, New York, NY, USA, 778–780. https://doi.org/10.1145/3462244.3479880

[Blog post discussing LLMs and provenance] Cargnelutti, M. (2023). *"Did chatgpt really say that?": Provenance in the age of Generative AI.: Library innovation lab*. RSS. https://lil.law.harvard.edu/blog/2023/05/22/provenance-in-the-age-of-generative-ai/

Engineering Ecosystems with AI

Alex Pentland, MIT

Federal Infrastructure bill includes social structure in definition of infrastructure

Why? Failures caused by not accounting for human behavior:

* **Pandemic and non-pharm interventions**
* **Finance (SVB, 2008, etc): 26 standard deviation events??**
* **Inequality, climate change action**

A simple example: social structure via "extended census" real-time flows of people

where the community works, shops, plays
minimal privacy, security risks; widely understood

health example: human flows and infectious disease

1% increase in vax rate yields

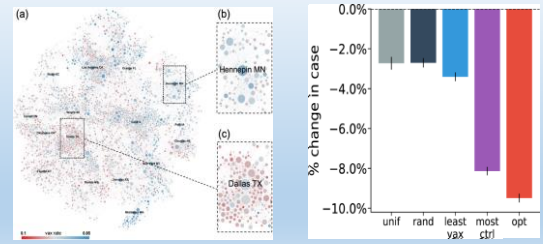

Implications of COVID-19 vaccination heterogeneity in mobility networks, Yuan, Jahani, Zhao, Ahn & Pentland, *Communications Physics* **volume 6**, Article number: 206 (2023)

wealth creation example

social exploration **is defined as diffusion of ideas and skills between communities, and is a major causal factor in:**

- **Spread of commerce between communities**
- **Innovation rates for US states in 1800's US**
- **Growth in commercial capabilities of countries**

and so predicts GDP growth in cities / nbhds around the world

wealth creation, social exploration, and transportation planning

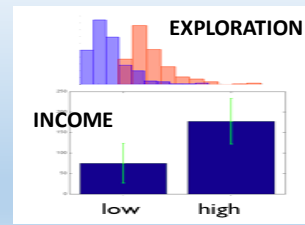

holds everywhere:
Economic outcomes predicted by diversity in cities, Chong, Bahrami, Chen, Balcisoy, Bozkaya, Pentland EPJ Data Science **9 (1), 17**

The predictability of consumer visitation patterns, Krumme et al. *Nature Scientific Reports* **3**, 1645 (2013)

and COVID made it worse….

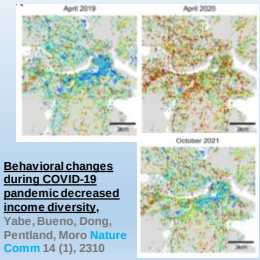

Behavioral changes during COVID-19 pandemic decreased income diversity, Yabe, Bueno, Dong, Pentland, Moro Nature Comm 14 (1), 2310

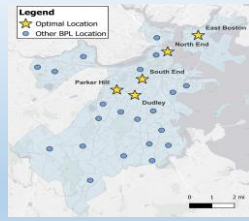

Nick Caros, MIT PhD thesis, 2023

can we do better than just relying on interaction?

we certainly need to!

O*NET job Transition data

Isabella Loaiza PhD Thesis, 2023

Hopeful signs:

Many recent studies are showing that Large Language Models (LLMs) like ChatGPT help

medium skill workers
more than
high skill workers

Why this might happen (human studies):

best action (posterior)= your estimate (likelihood) * popularity of action (prior)

AI can suggest actions you didn't think of

Bayesian Collective Learning Emerges from Heuristic Social Learning, Krafft,..,Pentland *Cognition* 2021

Accuracy-Risk Trade-off in …Financial Predictions, Adjodah,…Pentland, *Entropy* 2021

Kernel Methods for Cooperative Multi-Agent Contextual Bandits Dubey..Pentland (ICML), 2020

So:

build AI that augments human social exploration

best action (posterior)= your estimate (likelihood) * popularity of action (prior)

Estimate from on-line conversations
aka "common sense" -- Minsky

(be sure to ask: common to what community?)

Digital Conversation Space

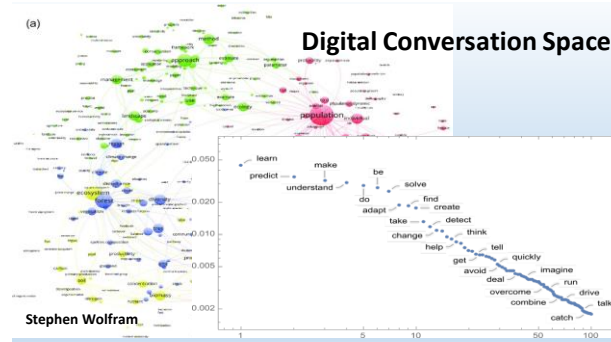

Stephen Wolfram

Repeat until you hit "end" token…

{ The best thing about AI is its ability to,
The best thing about AI is its ability to create,
The best thing about AI is its ability to create worlds,
The best thing about AI is its ability to create worlds that,
The best thing about AI is its ability to create worlds that are,
The best thing about AI is its ability to create worlds that are both,
The best thing about AI is its ability to create worlds that are both exciting,
The best thing about AI is its ability to create worlds that are both exciting, }

Repeat with noise to estimate **P(A|B)**

PERSONAL INVESTING

AI tools that aid exploration

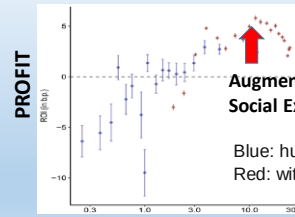

BUSINESS IDEAS

novelty
human vs GPT-4 ideas

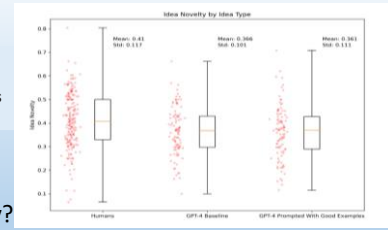

Proof of creativity?
no… AI tools that aid social exploration

Ideas Are Dimes A Dozen Girotra, Meincke, Terwiesch, Ulrich, *Wharton TR*

SCIENCE: citation "language" enables social exploration

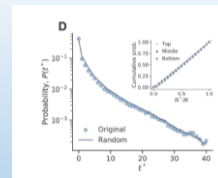

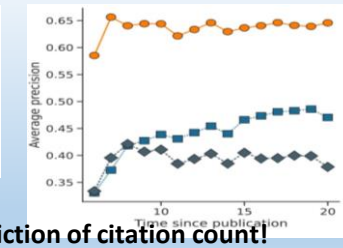

accurate prediction of citation count!
works with LAW and PATENTS too!

Kojaku, Mahari, Lera, Moro, Pentland, Ahn Nature in review

RNA "language" enables exploration of microbiome

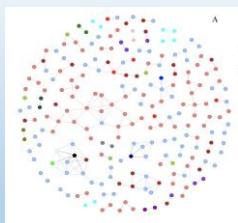

METHA.AI

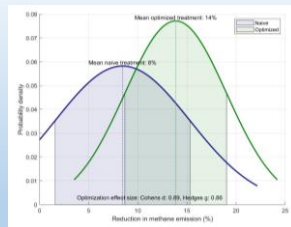

Naïve deployment (Agolin) vs. Optimized Use

Take home messages:

- Improve social exploration between communities using knowledge of social interaction patterns
- AI for enhancing social exploration may be a safe and effective way address many serious social problems

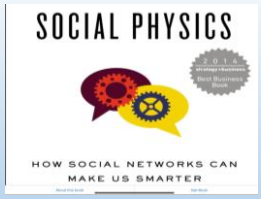

Prof. Alex Pentland, MIT
http://connection.mit.edu