Original Article

Community Detection in Medicine: Preserved Ejection Fraction Heart Failure (HFpEF)

Michael Liebman[1]* [iD], Stefania Pieroni[2], Michela Franchini[2], Loredana Fortunato[2], Marco Scalese[2], Sabrina Molinaro[2], Mark Wainger[3] and Steven P. Reinhardt[3]

[1]*IPQ Analytics, Kennett Square, Pennsylvania, United States;* [2]*Epidemiology and Health Research Lab, Institute of Clinical Physiology, National Research Council of Italy, Pisa, Italy;* [3]*Quantum Computing Inc., Leesburg, Virginia, United States*

Abstract

Background and objectives: The COVID-19 pandemic raised awareness of the complexities of the patient, the disease, and the practice of medicine. The impact of these reaches beyond healthcare (*e.g.*, supply chains, politics, socioeconomic factors) to include nations, individuals, and molecules. In personalized medicine, "accurate diagnosis" is critical as it affects patient management, clinical trial recruitment, regulatory approval, and reimbursement policies for payers. Conventional statistics evaluate hypothesis-driven reductionist practices in medicine, *e.g.*, the use of "scores" combining individual measurements, and are often limited by the data:variables ratio. True personalization (N of 1) is not practical but better stratification of diseases and patients can improve diagnoses. This work describes our approach and tests its ability to identify patient complexity and clinical markers in the trial of a candidate HFpEF drug better than prior methods.

Methods: This study evaluated *discovery* or *data-driven* approaches, by applying community detection (CD), forgoing statistical significance to identify unknown relationships. We reanalyzed data from the I-PRESERVE study of heart failure with preserved-ejection fraction, where subgroup analysis was unsuccessful. We initially performed unipartite CD analysis and evolved to address the complexity in real-world data using a bipartite model. The mathematically grounded modularity metric enabled greater confidence in community assignments.

Results: This reanalysis with CD revealed novel patient subgroups with stronger supporting rationale for group assignments, pointing to further refined and targeted studies.

Conclusions: We believe that generalization of the CD approach to higher-dimensional data can lead to a "next generation of phenotyping" that encompasses the temporal progression of the patient.

Introduction

The COVID-19 pandemic has raised awareness of the complex nature of the disease's impacts beyond healthcare to global networks (*e.g.*, economies, supply chains, politics, socioeconomic factors) in a hierarchy that ranges from nations to individuals to molecules (*e.g.*, vaccines, antibodies, viruses). Humans tend to deal with such complexity by applying reductionistic approaches to cut the problem into pieces that can be better conceptualized and managed. While this makes the approach more tractable, it can limit the ability for solutions to generalize to real-world problems. The need to rapidly digest, evaluate and create policy/recommendations based
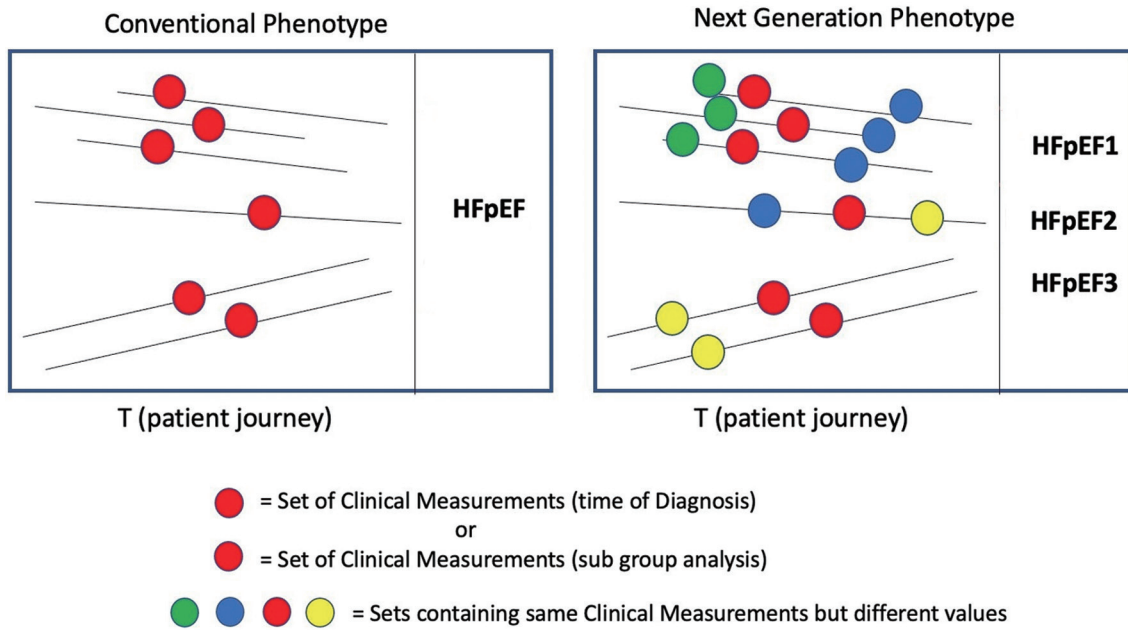
Fig. 1. Defining Next-generation Phenotyping for Disease Stratification.

on the increasing amount of data being generated in COVID-19 studies is constantly challenged by the lack of specificity that results from reductionist labeling. The seemingly simple classification of *fully vaccinated*, *partially vaccinated,* and *non-vaccinated* typifies this issue.

Algorithmic modeling has developed rapidly and can be used both on large complex data sets and on smaller, more quantitative data sets.[1] With the current emphasis on the capture and analysis of *big data*, one of the great challenges is the ability to compare or integrate diverse data types.[2] In this study, we choose a network-community-based approach that is founded on an algorithmic model. We implemented and have been exploring the use of community detection algorithms[3–5] that can be applied in two ways: (1) identify a target outcome and determine what factors are associated with predicting whether a population will attain that outcome; and (2) identify communities with common characteristics and evaluate their respective outcomes to facilitate better patient management, drug development, and more effective reimbursement policies. In Figure 1, we highlight the difference between these two approaches in the example of heart failure with preserved ejection fraction (HFpEF).

In the United States, heart failure affects approximately 6.2 million individuals, with a prevalence of 2.4–2.6%, and appeared on ~14% of all death certificates in 2018.[6] Heart failure is considered to be a "complex clinical syndrome" characterized by high comorbidity burdens.[7] Many of the patients exhibit non-specific symptoms, which makes it difficult to identify heart failure and distinguish it from other conditions. Thus, many patients may have undiagnosed heart failure, or even when diagnosed, other undiagnosed concomitant conditions, such as diabetes which is common in patients with acute heart failure, may confound the heart failure diagnosis. It is important to identify these patients and provide access to appropriate treatment to reduce mortality, improve healthcare, and reduce costs derived from undiagnosed/misdiagnosed diseases.[8] Partly due to this difficulty of a clear diagnosis, there currently are no drugs approved for use for HFpEF. The diagnosis and management of HFpEF remain challenging for the physician,

drug developers, payers, and ultimately for the patient.

The National Heart and Lung Institute (NHLI) Working Group on Research Priorities for HFpEF identified deep phenotyping as a critical need to address real-world complexity.[9] Our thesis is that community detection methods may support HFpEF risk stratification, which would be doubly promising because they are expected to be greatly accelerated by early quantum computers. This work aims to describe our approach and test its ability to identify patient complexity levels and related clinical markers in the trial of a candidate HFpEF drug better than prior methods.

Methods and materials

Study design

This study re-analyzes the data from the Irbesartan in Patients with Heart Failure and Preserved Ejection Fraction (I-PRESERVE) clinical trial (where the initial analysis detected no subpopulation benefiting from the drug), recognizing that the graph representing the patients and medical variables is inherently bipartite and hence the community detection performed on it should reflect that bipartiteness. The data consists of 11 medical variables with a total of 34 categorical values (see Supplementary File 1). The community-detection implementation is also believed to provide better answers due to higher compute intensity, with the added benefit of further improved effectiveness from early quantum computers.

Data

This study was carried out using the baseline data derived from the Irbesartan in Patients with Heart Failure and Preserved Ejection Fraction (I-PRESERVE) clinical trial involving more than 5,000 patients, which began in 2002 and extended over 5 years.[10,11] The data were obtained by contacting the I-PRESERVE study authors.

Table 1. Inclusion/Exclusion Criteria for I-PRESERVE

Inclusion Criteria	Exclusion Criteria
a. At least 60 years of age	a. Previous intolerance to an angiotensin-receptor blocker
b. Heart failure symptoms	b. Alternative probable cause of the patient's symptoms (*e.g.*, significant pulmonary disease)
c. Left ventricular ejection fraction of at least 45%	c. Any previous left ventricular ejection fraction below 40%
d. Required patients to have been hospitalized for heart failure during the previous 6 months	d. History of the acute coronary syndrome, coronary revascularization, or stroke within the previous 3 months
e. Current New York Heart Association (NYHA) class II, III, or IV symptoms with corroborative evidence	e. Substantial valvular abnormalities
f. If they had not been hospitalized, they were required to have ongoing class III or IV symptoms with corroborative evidence, *e.g.*	f. Hypertrophic or restrictive cardiomyopathy
i. Pulmonary congestion on radiography	g. Pericardial disease; cor pulmonale or other cause of isolated right heart failure
ii. Left ventricular hypertrophy or left atrial enlargement on echocardiography	h. Systolic blood pressure of less than 100 mm Hg or more than 160 mm Hg
iii. Left ventricular hypertrophy or left bundle-branch block on electrocardiography	i. Diastolic blood pressure of more than 95 mm Hg despite antihypertensive therapy
g. Treatment with an angiotensin-converting enzyme (ACE) inhibitor was permitted only when such therapy was considered essential for an indication other than uncomplicated hypertension	j. Other systemic diseases limit life expectancy to less than 3 years
	k. Substantial laboratory abnormalities (such as a hemoglobin level of less than 11 g per deciliter, creatinine level of more than 2.5 mg per deciliter [221 Î mol per liter], or liver-function abnormalities)
	l. Characteristics that might interfere with compliance with the study protocol

Information from www.clinicaltrials.gov.

Initial trial results showed no benefit over placebo and subsequent multiple subgroup analyses were attempted using traditional statistical clustering approaches with minimal success.[12]

Clinical guidelines for heart failure

The development of guidelines, commonly by committees based on data from randomized clinical trials, typically reveals limitations both in the assignment of specific diagnoses and their subsequent use in determining appropriate treatment. The development of guidelines for a specific condition ideally includes the recognition of the real-world complexity of the patient and disease with the need to differentiate accurately among both disease and patient sub-groups.

The challenge in the diagnosis of preserved-ejection-fraction heart failure (HFpEF) reflects that of applying clinical guidelines to address a syndromic condition. A gap exists between current clinical practice applied to real-world patients and strict adherence to either European Society for Cardiology (ESC)[13] or AHA[14] guidelines, which themselves undergo independent, periodic updating. It should be noted that such guidelines are intended to provide informed guidance to clinicians and full compliance is not mandated (nor expected). For example, the threshold level for LVEF as *preserved* varies among groups and reflects the observation, *e.g.*, in the Framingham study, of limited ability to assign a specific

threshold, especially because of the potential that an individual patient's value may have changed due to pre-treatment between the time of initial diagnosis and enrollment into the I-PRESERVE trial and the more general observation that some patients with LVEF <45% may have HFpEF.[15]

Inclusion/exclusion criteria for I-PRESERVE

The inclusion and exclusion criteria for the I-PRESERVE study can be found at www.clinicaltrials.gov and was listed in Table 1.

Sparseness of real-world (Clinical Trial) data

The current trend/discussion in data analysis, in healthcare and many other domains, focuses on access to and analysis of big data, but it has been long known that there is a constant tension between quantity and quality of data.[16] Many current analytic methods, *e.g.*, machine learning and deep learning, are dependent on access to large data sets; this reflects their emphasis on correlative vs causal analysis. For many applications, correlative analysis can provide critical guidance and optimal results but in medicine, unknown biases that may be present in the data may limit the utility of such analyses and even result in incorrect results and interpretation. The reality of real-world clinical data is its sparseness, *i.e.*, meas-

urement of limited numbers of medical variables and rarely in a continuous manner over time. The anticipated transition to digital medicine will help address this issue but will require significant evolution of clinical practice, physician compliance, and patient adherence so will develop slowly over time despite increasing access to technology. Analytic methods, therefore, will be confronted with sparse data sets for some time and need to be pragmatic in their approach. Clinical trials provide a more controlled environment for the collection of data than typical clinical data, *e.g.*, electronic health records (EHR), because of requirements to follow specific protocols, but even these exhibits significant sparseness in data collection. Table 2 documents the number of patients for whom data was gathered, at each time point in the I-PRESERVE data. We observe, *i.e.*, at Month 18, that for almost all measurements more than half the patients do not have data. In general, high-density data collection is expensive and typically not undertaken without the ability to show value for the effort. In this study, we show how community detection can show increasing value for data integration, even in an incremental manner.

The data presented in Table 2 reflect that collected in I-PRESERVE based on the protocol and case report forms. One goal of the analysis is to enable ease of integration of the results into current clinical practice. To facilitate this, the data was further mapped into conventional clinical panels used in the diagnosis and patient management. Note that this results in some observations being present in more than one panel, *e.g.*, Alanine Aminotransferase (ALT, also known as Serum Glutamic Pyruvic Transaminase, SGPT) as shown in Table 3.

Current practice

Humans tend to deal with the complexity of many real-world data by applying reductionistic approaches to cut a problem into pieces that can be better conceptualized and managed.[17] While this makes the approach more tractable, it can limit the ability for solutions to generalize to real-world problems. The application of reductionism in biology has been shown to be self-limiting.[17,18] This is of particular concern when applied to diseases and disease management as the limitations of "naming" (classifying) a condition can have a significant impact on treatment decisions, payer reimbursement, and drug development, all resulting in sub-optimal patient management.

The power of algorithmic modeling has caused researchers to want to combine or integrate diverse types of data, and the novelty of these combinations has further led to a desire to examine data using different algorithms. For example, for data sets that do not necessarily present readily definable clusters, the application of different clustering methodologies may result in variable results which may make any interpretation dependent upon the methodology used.[2]

In the United States, heart failure affects approximately 6.2 million individuals, with a prevalence of 2.4–2.6%, and appeared on ~14% of all death certificates in 2018. Globally it is estimated that 64.3 million people are living with heart failure or ~1–2% of the general population. In the US, the cost of care for heart failure, including direct and indirect costs, is estimated at $43.6B per year and projected to increase to $69.7B by 2030 with ~70% of these costs going to medical care.[6]

Heart failure is commonly classified in terms of the Left Ventricular Ejection Fraction (LVEF) into three classes: heart failure with *reduced* (HFrEF; LVEF <40%, previously known as systolic heart failure), *mid-range* (HFmEF; LVEF 40–49%), or *preserved* ejection fraction (HFpEF; previously known as diastolic heart fail-

ure, LVEF ≥50%).[19] These thresholds may vary among studies and sometimes mid-range is further divided into 40–45% and 45–50%. The actual observed distribution reveals the challenge in defining separable boundaries using only LVEF as the major classifier (or label). The data in Table 4[20] display the association between simple LVEF classifications, gender, and the causes of death from cardiovascular diseases (CVD), distributed into the coronary heart (CHD) and other diseases.

Actual diagnostic guidelines, however, include additional factors and clinical/medical variables, *e.g.*, comorbidities and levels of N-terminal pro-B-type natriuretic peptide, to establish the diagnosis and highlight the complexity of disease presentation. As noted above, the reductionist classification of the disease, based solely on LVEF, does not adequately stratify the disease and patients and hence enables more personalized diagnosis and management and/or development of more effective drugs.

Patients who currently present or have prior symptoms of heart failure are classified as HFpEF. The American College of Cardiology (ACC)/American Heart Association (AHA) classifies these patients in stages C and D, while those patients in stage B are considered to be at risk for developing HFpEF. Additionally, HFpEF must be distinguished from valvular disease, pericardial disease, and cardiac amyloidosis. Currently, approximately 50% of heart failure is HFpEF with a higher prevalence among older patients and females. Moreover, HFpEF diagnosis has increased by 45% over the last two decades.

The other 50% of heart failure is classified as HFrEF. Similar clinical manifestations appear in HFrEF and HFpEF including peak oxygen uptake (VO2) and neurohumoral activation. Many comorbidities are common between HFrEF and HFpEF including hypertension, atrial fibrillation, diabetes mellitus, metabolic syndrome, obesity, chronic obstructive pulmonary disease (COPD), chronic kidney disease, and anemia.

Angiotensin-converting enzyme inhibitors (ACEIs), angiotensin receptor blockers (ARBs), beta-blockers, mineralocorticoid receptor antagonists (MRAs), and diuretics form the basis of first-line pharmacological management of left ventricular heart failure with reduced ejection fraction (*i.e.*, HFrEF). However, until 2021, no drugs had been approved for use in HFpEF although there were 17 active clinical trials involving 14 unique agents and testing 14 endpoints, involving 10 distinct classes of mechanism of action, and thus there was great interest in finding an effective drug. In 2021, based on analysis of the PARAGON-HF trial,[21] sacubitril/valsartan (Entresto™, Novartis) received a broad heart failure indication that reached into the normal range of ejection fraction. It was noted that most benefits remained in the HFrEF population despite missing its primary endpoint. Significance was shown in subgroup analysis involving patients with an ejection fraction at or below the median of 57%. (Note: More recently, a sodium-glucose cotransporter-2 (SGLT2) inhibitor (Farxiga™) has shown a positive effect in HFpEF patients in the Emperor-Preserved study.[22])

The main tests that comprise the initial HFpEF diagnosis remain Doppler echocardiography and serum natriuretic peptide levels. Further diagnostic scoring of patients currently utilizes two scores, H2FPEF and HFA-PEFF, which include some degree of subjectivity in the evaluation and interpretation.[23] H2FPEF includes evaluation of body mass index, hypertension, atrial fibrillation, pulmonary hypertension, age, and filling pressure. HFA-PEFF incorporates assessment of major and minor criteria within functional, morphological, and biomarker categories. In general, however, the use of multi-variable scores can obscure critical heterogeneity in patient groups. It has been noted that current HFpEF diagnoses are confounded by the presence of several significant subtypes.

Table 2. Data collected for I-PRESERVE at baseline, intermediate, and final times, reflecting the sparsity of current clinical practice

General	Specific	Exam	Base-line	Wee 2	Wee 8	Mon 6	Mon 10	Mon 14	Mont 18	Mon 30	Mon 42	Mon 54	Mon 66	Tot
Liver & Kidney Function	Liver function tests	Alanine Aminotransferase (ALT)	4,302	31	31	24	19	15	18	21	15	20	4	2,102
Liver & Kidney Function	Liver function tests	Aspartate Aminotransferase (AST)	4,302	32	32	22	19	16	22	26	16	21	4	2,104
Liver & Kidney Function	Liver function tests	Bilirubin, Total	4,302	30	27	17	14	10	9	12	10	16	2	2,104
Liver & Kidney Function	Kidney function tests	Blood Urea Nitrogen	4,209	27	27	18	14	10	9	9	10	16	2	2,058
Liver & Kidney Function	Kidney function tests	Creatinine	4,361	3,811	3,703	3,552	72	71	2,840	2,563	1,890	915	48	2,105
Liver & Kidney Function	Kidney function tests	Creatinine Clearance (MDRD)	4,361	3,811	3,699	3,549	71	70	2,834	2,557	1,888	914	46	2,105
Other Chemistry Tests	Protein tests	Albumin	4,302	30	27	17	14	10	9	11	10	16	2	2,103
Hematology I	Erythroc./Platel. attributes	Hematocrit	4,153	24	16	10	10	6	6	6	10	16	2	2,072
Hematology I	Erythroc./Platel. attributes	Hemoglobin	4,153	24	16	11	13	6	8	11	10	16	2	2,072
Hematology I	Erythroc./Platel. attributes	Platelet Count	4,146	24	15	9	11	5	4	6	11	16	2	2,068
Hematology II	Quantitative WBC	Leukocytes	4,152	24	16	11	11	6	6	6	11	17	2	2,072
Hematology II	WBC differential count	Neutrophils (absolute)	4,125	24	15	9	9	6	4	6	10	13	2	2,064
Blood	Other testing	NT-proBNP	3,620	32	13	3,034	1,160	2,926	89	7	4	0	0	0
Electrolytes	Electrolytes	Potassium, Serum	4,316	3,805	3,691	3,493	72	68	2,822	2,559	1,887	915	47	2,093
Electrolytes	Electrolytes	Sodium, Serum	4,302	29	27	17	14	11	9	12	10	17	2	2,103

Table 3. Medical variables were tracked in the I-PRESERVE study, by the system

Blood	Liver	Kidney	Spleen
Age	Age	Age	Age
Gender	Gender	Gender	Gender
EOS – Eosinophils	ALB – Albumin	ALB – Albumin	ALT – Alanine Aminotransferase (SGPT)
HCT – Hematocrit	ALP – Alkaline Phosphatase	ALP – Alkaline Phosphatase	BILI – Bilirubin
HGB – Hemoglobin	ALT – Alanine Aminotransferase (SGPT)	BICARB – Bicarbonate	HCT – Hematocrit
LYM – Lymphocytes	AST – Aspartase Aminotransferase (GOT)	CL – Chloride	HGB – Hemoglobin
MONO – Monocytes	BILI – Total Bilirubin	CREAT – Creatinine	PLAT – Platelets
PLAT – Platelets	GGT – Gamma Glutamyl Transferase	K – Potassium	RBC – Red Blood Cell/Erythrocytes
RBC – Red Blood Cells/Erythrocytes	GLUC – Glucose	Na – Sodium	SPLEENLEN = Numeric Spleen Length
NEUT – Neutrophils		UR - Urate	

Table 4. Distribution of left-ventricular-ejection-fraction classifications, gender, and disease

	Cardiovascular Disease (CVD) Deaths			
LVEF classification by gender	CHD	Stroke	Other CVD	Total
HFrEF male	45%	5%	27%	77%
HFrEF female	30%	14%	26%	70%
HFpEF male	11%	3%	25%	39%
HFpEF female	15%	11%	23%	49%

Modularity-based community detection

The goal of this study was to classify or stratify patients using community detection algorithms that were objectively data-driven, *i.e.*, which identified patient groups based on similarity of clinical presentation. This was done differently from conventional subgroup analysis that would select a target characteristic, *e.g.*, response to a specific therapy, and then identify the characteristics that were common among those patients. In addition, the community detection method requires no pre-determination of how many patient groups, how many medical variables were needed to define these groups, whether each group reflected different values of the same medical variables, or even if the same set of variables was used to define the individual groups based on their values. The community detection algorithms were evaluated using both unipartite and bipartite graphs.

Modularity-based community detection (mCD) was first described and implemented by Newman and Girvan[3] based on their insight that communities in a graph are best defined as "a statistically surprising arrangement of edges". Their analysis converted the general problem of finding communities into a graph-based constrained optimization problem, where the metric to be maximized is *modularity*, and there are constraints for every node to be in exactly one community. Modularity is defined as a difference of two terms: (1) the density of edges inside communities as compared to edges between them, minus (2) the same measure for the corresponding *null model*, *i.e.*, a graph where each node has the same number of edges as the original graph, but the connected nodes are randomized.

Finding the globally optimal answer to a modularity maximization problem is *NP-hard*, meaning its computational cost on a classical computer grows exponentially with the number of nodes. Thus, many implementations, including Newman-Girvan's, are *greedy* heuristic algorithms that make locally optimal decisions each iteration, with no guarantee they will be able to find the globally optimal solution. Quantum computers, whose quantum bits (*qubits*) work in the exponentially larger quantum problem space than *bits* of classical computers, are expected to be able to find globally optimal answers for many classes of NP-hard problems efficiently, and there is vigorous research into quantum algorithms even though practical hardware implementations for real-world sized problems are still years away.[24] Researchers at Los Alamos National Laboratory (LANL) describe a quantum implementation for mCD in Negre *et al.*[25] that targets the globally optimal answer, though real-world samplers are currently heuristic. If there are N nodes in the graph, searching for K communities requires N * K variables, or qubits if the problem is being solved on a quantum computer. In the case of the current study, N = (3,935 patients + 34 medical variables) and K = 5, which requires ~20,000 densely connected qubits. This is well beyond the capability of the largest available quantum annealing computer, the D-Wave Advantage™ system, with 5,600 sparsely connected qubits, and even further beyond the capability of the largest available gate-model quantum

computer, the IBM Eagle processor, with 127 sparsely connected qubits. The effectiveness of early quantum computers in solving constrained optimization problems like mCD was discussed by Hen and Spedalieri[26] and Hadfield *et al.*[27]

Note that our hypothesis consists of two primary tenets: that mCD will give better clustering than prior methods, and that quantum computers will accelerate the performance of mCD. While we provide conceptual arguments supporting the latter, only the first of these is tested in this paper.

Community detection for unipartite and bipartite graphs

The simplest graphs are *unipartite*, *i.e.*, they consist of only one type of node. An example would be to consider the atoms in a protein molecule as the nodes of a graph, with the strength of their connections equal to the pairwise atomic-level forces between them. mCD can identify a community of atoms for each amino acid in the protein. Unipartite graphs were what Newman and Girvan originally studied, and many software systems only consider this type of graph.

However, in the real world, many graphs are *bipartite* graphs, which are defined as having nodes of two types and edges that only join opposite types of nodes. Communities for unipartite graphs are often described informally as having high connectivity within communities and lower connectivity between communities. That mental model does not hold for bipartite graphs, where there is, by definition, no connectivity between same-type nodes within a community, so we must depend more explicitly on the definition of modularity. To bring the bipartite model to a concrete example, we view humans, not as atoms; they do not generate their own connections. They are connected by the events they attend, the papers they co-author, the movies they act in, *etc.*, and so they are typically found in bipartite graphs. One standard example studied in the literature is the Southern women graph documented by Davis *et al.*,[28] consisting of 18 women and the 14 events they attended in the 1930s. The women are connected to the events they attended, and the events are connected to the women that attended them. Figure 2 illustrates the best assignment of communities found by Liu and Murata[29] for the Southern women graph. Each community (to the left and right of the vertical blue bar, respectively) contains both events (white nodes with black text) and women (black nodes with white text).

In the current study, there are two types of nodes – patients and medical variables – so it is also a bipartite graph. Medical variables do not connect directly to other medical variables and patients do not connect directly to other patients; they only connect indirectly through common medical data.

The null model for a unipartite graph is not correct for a bipartite graph, because it assumes any two nodes can be randomly connected by an edge, and so it connects nodes of the same bipartite type in violation of the definition. Barber[30] presents the correct null model for a bipartite graph, where nodes are randomly connected only to nodes of the opposite type. See Calderer and
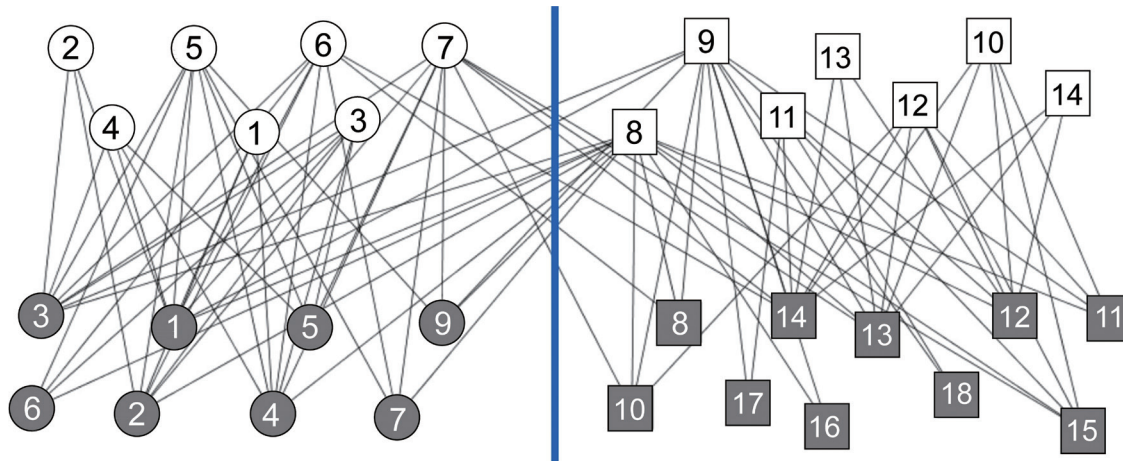
Fig. 2. Liu and Murata community assignment for Southern women bipartite graph.

Kuijjer[31] and Ganji[32] for more discussion of when unipartite or bipartite mCD is appropriate.

The current study did the initial mCD analysis[33,34] using the Gephi implementation of the Louvain method.[35] Gephi is an opensource graph visualization tool,[36] which uses a heuristic algorithm that is limited to the computational resources present where it is executing, usually a user's laptop/desktop system. That limitation directly affects the quality of the community assignments it can find. Gephi calculates modularity only for unipartite graphs. We calculated bipartite modularity for the Gephi unipartite solution by running it through Qatalyst's bipartite modularity calculation.

The Qatalyst quantum-acceleration platform, by Quantum Computing Inc., samples binary constrained optimization problems using classical and quantum processors, with quantum-ready heuristic formulations; the best results are currently obtained running purely classically, with no quantum contribution.[37,38] Graphbased mCD is readily expressed as a constrained optimization problem[25] that Qatalyst can effectively sample. Qatalyst runs on AWS servers, with the compute-intense classical quadratic-unconstrained-binary-optimization (QUBO) sampler executing on thousands of cores. The current study used Qatalyst for both unipartite and bipartite calculations.

Selection of data for analysis: framing the question

Understanding the complexities of data like those present in this study has led to much development and application of methods such as deep learning. Our focus on moving from correlative towards causal analysis and the ability to calculate real-world results has led us to enable the evaluation of specific models that are readily applied in current clinical practice. Several example models include:

1. Disease Model 1 (Patient Demographics and anamnesis): Age, Gender, BMI, Age at Diagnosis, Number of years post HF diagnosis (entry into the trial), Atrial Fibrillation by ECG, Left Bundle branch block by ECG, Left Ventricular Hypertrophy by ECG, Peripheral Edema, Left Ventricular Ejection Fraction, Etiology;
2. Disease Model 2 (Clinical History): Age, Gender, BMI, Age at Diagnosis, Number of years post HF diagnosis (entry into the trial), History of COPD, History of Diabetes, History of Atrial Fibrillation, Heart Failure within previous 6 months, Jugular Venous Distension, Lung Sounds, Left ventricular hypertrophy or Left Atrial Enlargement, NY Heart Association Functional

Classification;
3. Hematologic Profile (Clinical Data): Age, Gender, BMI, Albumin, Hematocrit, Hemoglobin, Platelet Count, Leukocytes, Neutrophils (absolute), NT-proBNP;
4. Liver & Kidney Function. (Clinical Data): Age, Gender, BMI, ALT, Aspartate Aminotransferase (AST), Bilirubin (total), Blood Urea Nitrogen (BUN), Creatinine, Serum Potassium, Serum Sodium, Creatinine Clearance (MDRD);
5. NT-proBNP (Clinical Data): At the time of initiation of I-PRESERVE, levels of NT-proBNP were not incorporated into clinical guidelines for the diagnosis of heart failure but were added in subsequent studies and are currently used as a threshold for diagnosis of heart failure;
6. Longitudinal/Temporal analysis (Clinical Data): Initial analysis of patient progress during the study was planned to develop patient trajectories, *e.g.*, patterns of progression both with treatment and placebo, for purposes of comparative analysis. Longitudinal/Temporal analysis was limited by data sparsity.

We choose to define a single model for liver function and kidney function, based on evidence for co-existing liver and kidney pathology in patients with chronic liver disease. Chronic liver disease is associated with primary and secondary kidney disease and impacts markedly on survival.[39] Moreover, we define the Hematologic model including NT-proBNP data as most HFpEF patients have elevated NT-proBNP levels. The NT-proBNP concentrations were related to baseline characteristics generally associated with worse outcomes for HF patients.[40]

For example, the clinical data for both the Liver & Kidney panel and the Hematologic panel are provided (in Supporting Material) where the categorization was based on observed medical-variable ranges in patients and also includes gender-based differences. We initially developed categorical boundaries, *i.e.,* cutpoints, for each medical variable based on current laboratory standards. These boundaries were further refined based on cardiologist input as potentially relevant to the study population. These boundaries also reflect the expected differences between male and female patients, and where appropriate, reflected a high/normal/low classification. A result is several categories that defined individual nodes: for example, BMI is defined into 5 categories Underweight (<18.5; BMI-L), Normal weight (18.5–24.9; BMI-N), Overweight (25.0–29.9; BMI-H-OV), Obese (30.0–34.9; BMI-H-OB), Morbidly obese (>35; BMI-H-OB*); ALT is defined by 3 categories for male (H high, N normal, L low), ALT-L <0, ALT-N between 0 and 55,
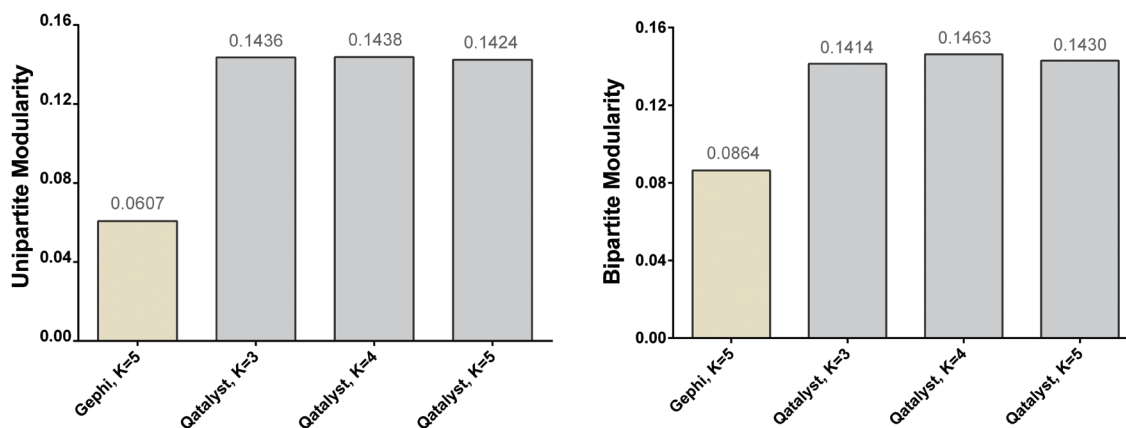
Fig. 3. Analysis of left-ventricular-ejection-fraction (LVEF) data, both for unipartite and bipartite, comparing GEPHI and Qatalyst.

ALT-H: >55 and 3 categories for female, ALT-L <0, ALT-N between 0 and 40, ALT-H >40.

Benefits of community detection for this analysis

Despite some observations being present in more than one panel, *e.g.*, ALT, the data was readily incorporated into this community-detection analysis.

We believe that the use of the community detection method described in this report can effectively address critical issues in clinical medicine, going beyond correlation to approach causality. Perhaps the leading example of these issues is that the use of the panels outlined in Table 3 provides a convenient assessment of a patient's status along with specific pathophysiologic domains through the highlighting of "outliers" from normal lab values:

1. The "normal" ranges for these medical variables may be dependent on an individual's clinical history, co-morbidities, diet, *etc.*, and thus require "personalized" evaluation;
2. While individual outliers may suggest diagnostic and therapeutic intervention, including lifestyle and/or medication, *e.g.*, low hemoglobin suggesting anemia, it is not uncommon for multiple medical variables to be non-normal with the increased complexity being less commonly observed and with reduced indications for management;
3. Temporal changes in an individual's multiple medical variables may be much more informative of a patient's status than single-point-in-time measurements. Such temporal patterns may involve clinical medical variables that never individually trigger an "abnormal" classification;
4. Higher level complexity in temporal measurements, *i.e.*, patterns involving more than one clinical variable, would be very difficult to detect but may be critical to define a more accurate diagnosis and staging of a specific condition.

Results

Unipartite analysis

Our first mCD-based analysis of the data viewed the problem as unipartite. The resulting communities via the Gephi implementation are described below in terms of disease characteristics' association and patient numerosity; the number of nodes always includes patients' and characteristics' nodes. With the unipartite Gephi implementation, the best results were obtained for K = 5 communities with modularity = 0.061 (Fig. 3).

Pink community

This group is composed of 1,125 nodes (1,116 patients). It aggregates Females aged 60–69, obese and morbidly obese (BMI-H-OB and BMI-H-OB*), showing several high serum values (Alanine Aminotransferase – ALT-H, Aspartate Aminotransferase – AST-H, Sodium – SOD-H) and normal values of Potassium (K-N) and Bilirubin (BILI-N).

Cyan community

It is the biggest community, composed of 1,177 nodes (1,168 patients). It aggregates Males, Overweight (BMI-H-OV), aged 70–79, with mildly reduced kidney function (KIDFUN-MIL) characterized by low values of BILI and K (BILI-L, K-L) and normal values of Blood Urea Nitrogen (BUN-N), Creatinine (CREAT-N), and SOD.

Green community

It is composed of 393 nodes (390 patients). It aggregates patients with normal weight (BMI-N) and low values of SOD (SOD-L), and normal values of AST (AST-N).

Red community

Composed of 999 nodes (990 patients), it aggregates the oldest patients (≥80), underweight (BMI-L) with severely or moderately reduced kidney function (KIDFUN-SEV, KIDFUN-MOD), and several abnormal values: high values of Bilirubin, Blood Urea Nitrogen, Creatinine, and Potassium (BILI-H, BUN-H, CREAT-H, K-H) and low values of Alanine Aminotransferase (AST-L), with normal values of Alanine Aminotransferase (ALT-N).

Yellow community

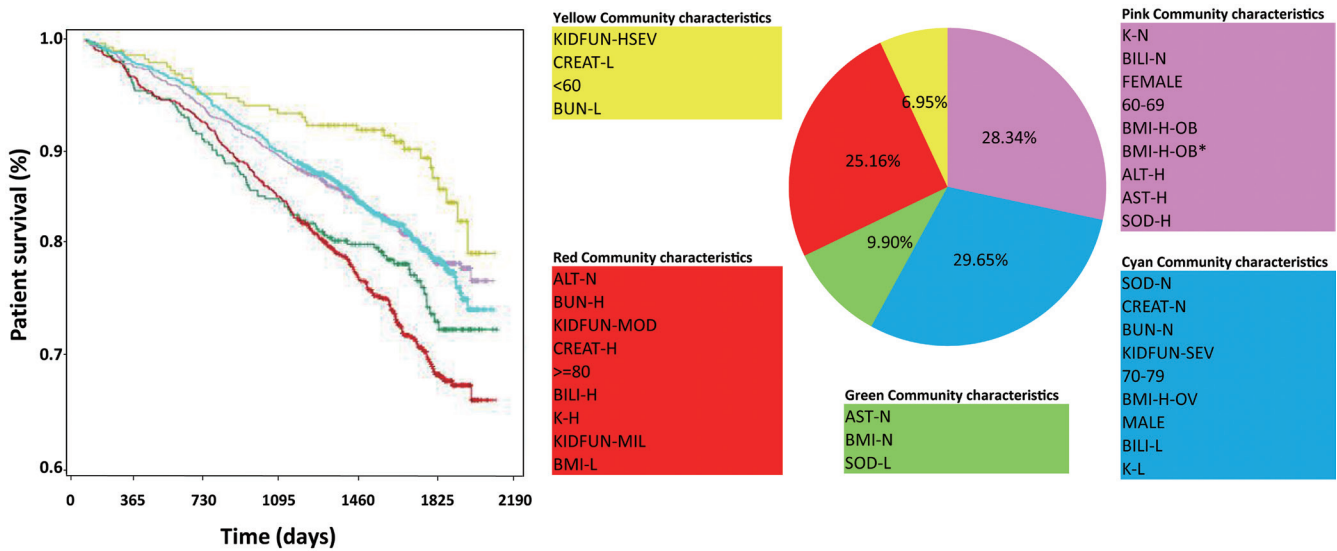This is the smallest community, composed of 276 nodes (272

Fig. 4. Communities found among Liver & kidney panel – Kaplan-Meier survival curve and defining medical variables. The variable abbreviation description can be found at Supplementary File 1.

patients). It aggregates the youngest patients (<60) with normal kidney function (KIDFUN-NOR), and low values of BUN and CREAT (BUN-L, CREAT-L).

To evaluate the clinical significance of these communities, we compared them using conventional survival analysis with the outcome of death/survival, neither of which had been included in the derivation of these communities. In terms of the quality of the patterns we can observe that the Red curve (representing the Red community) refers to the patients with shorter life expectancy; the Yellow curve (representing the Yellow community) shows the patients with higher life expectancy. This is coherent with the Liver & Kidney communities' composition. Moreover, the performed log-rank test (Bonferroni adjustment for multiple comparisons) revealed statistical significance: a) the Green curve has statistical significance concerning the Yellow ($p = 0.003$); b) the Red curve concerning the Yellow, Cyan, and Pink curves (Red vs Yellow $p < 0.001$; Red vs Cyan $p < 0.001$; Red vs Pink $p < 0.001$).

Using the same data, Qatalyst finds an optimal answer of K = 4 communities with modularity = 0.144 and finds a slightly smaller modularity (=0.142) for K = 5 communities (greater modularity indicates stronger communities). These values are graphically represented in Figure 3. Looking at the Kaplan-Meier survival curves in Figure 4, the lines for two of the five communities (Cyan and Pink) almost completely overlap, lending weight to the finding of four communities as being globally optimal.

Table 5 explores the overlap between the 5-community solu-

tions found by Gephi and Qatalyst. Each entry is a pair of numbers (m, p), where m is the number of medical variables and p is the number of patients in the community: so, by example, $(3,390)$ to the right of "Green (m, p)" indicates that in the Green community of the Gephi solution there are 3 medical variables and 390 patients. The outer values are the counts for the two solutions, horizontal values come from the Qatalyst solution and vertical values come from the Gephi solution. The values in the box represent the overlap of patients between pairs of communities. The assignment of numbers to communities is arbitrary, so even if the same solution was found twice, the community numbers assigned could be different. In this case, we see that a majority of Gephi Community 2 (the Red community) wound up as the majority of Qatalyst Community 5 (data in *italic*).

Bipartite analysis

Realizing that the data of the current study is inherently bipartite, with patients and medical variables comprising the two separate sets of entities, we performed the calculations again treating the graph as bipartite.

Qatalyst implements Barber's bipartite modularity calculation, which gives the unipartite Gephi solution bipartite modularity that is substantially higher than its unipartite modularity, a result due solely to the different null models. Maximizing bipartite modular-

Table 5. Comparison of unipartite communities found by Gephi and Qatalyst; (m, p), m = number of medical variables; p = number of patients

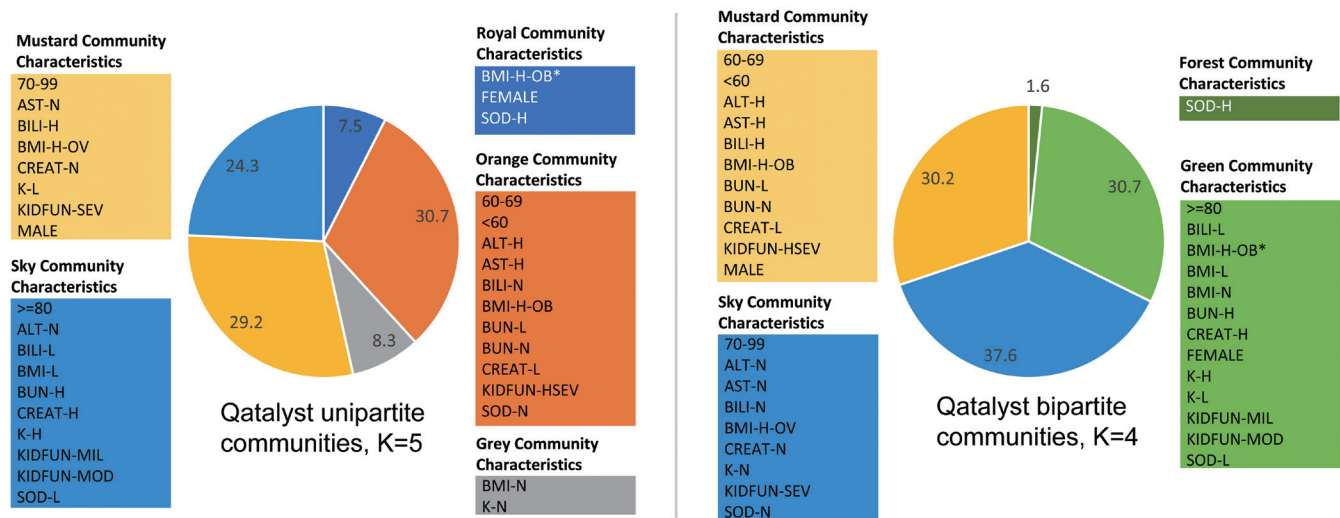|	Qatalyst Communities	Community 1 *(m, p)*	Community 2 *(m, p)*	Community 3 *(m, p)*	Community 4 *(m, p)*	Community 5 *(m, p)*
Gephi Communities		(3, 296)	(11, 1,208)	(2, 327)	(8, 1,149)	(10, 955)
Green *(m, p)*	(3, 390)	(0, 14)	(0, 59)	(1, 98)	(1, 137)	(1, 82)
Red *(m, p)*	(9, 990)	(0, 28)	(0, 148)	(0, 70)	(1, 228)	*(8, 515)*
Pink *(m, p)*	(9, 1,116)	(3, 182)	(5, 473)	(1, 78)	(0, 210)	(0, 173)
Yellow *(m, p)*	(4, 272)	(0, 5)	(4, 231)	(0, 8)	(0, 23)	(0, 5)
Cyan *(m, p)*	(9, 1,168)	(0, 67)	(2, 297)	(0, 73)	(6, 551)	(1, 180)

Fig. 5. Characteristics of communities found by Qatalyst, unipartite *vs.* bipartite. The variable abbreviation description can be found at Supplementary File 1.

ity, Qatalyst again finds a solution with notably higher modularity than Gephi, *i.e.*, modularity = 0.146 for K = 4 communities, as shown in Figure 5, with different communities than it found when ignoring the bipartite structure. For Qatalyst, the modularity difference between K = 4 and K = 3/K = 5 is larger using the bipartite null model than the unipartite, but the differences are still small.

Comparing the community assignments in detail between Qatalyst results for bipartite K = 4 and unipartite K = 5 in Table 6, we see a strong overlap between bipartite Community 1 and unipartite Community 5, as well as bipartite Community 3 and unipartite Community 2 (both in *italic*). Most of the small bipartite Community 4 winds up in unipartite Community 1, but bipartite Community 2 is redistributed across the unipartite communities.

Comparing the characteristics that describe the communities found between Qatalyst unipartite and bipartite approaches (see Fig. 5), we see that some communities have similar characteristics, such as the Orange unipartite and the Mustard bipartite community, but there are still some significant differences between the two. The difference in the obtained results demonstrates the importance of using the correct null model for the type of mCD problem being solved.

Discussion

In this study, we have applied a novel approach to stratification

of HFpEF patient data to objectively elucidate the HFpEF subtypes using a hypothesis-free, data-driven approach and outlined the method and results.

In the last few decades, various technological and cultural changes have contributed to exposing the limitations of the hypothesis-driven approach to knowledge discovery.[41] First, all scientific disciplines are nowadays required to tackle increasingly challenging, nonlinear problems and systems, some of which are very difficult, if not impossible, to model with theories based on first principles. Moreover, many newly interesting phenomena, for various reasons ranging from intrinsic randomness to inaccessibility for measurement, are characterized by a high level of uncertainty, limiting the effectiveness of traditional statistical approaches. All this is in the context of an exponential increase in data availability and complex relations among different data collected on the same statistical unit, in particular, if it is the real-world patient. The imprecision of many common diagnostic categories implies the need of specifying inclusion/exclusion criteria in more detail for a clinical trial, along with scientific and commercial considerations. This leads to significantly differentiating the patient recruited for the trial from the real-world patient and limiting the ability to directly compare results among independent trials.[42,43]

The large amounts of data being captured in COVID-19 studies and the urgency of processing that data into information usable by laypeople have highlighted the lack of specificity that results from reductionist labeling. The seemingly simple classification of

Table 6. Comparison of unipartite and bipartite communities found by Qatalyst; (m, p) m = number of medical variables; p = number of patients

Qatalyst	Bipartite	Community 1 (m, p)	Community 2 (m, p)	Community 3 (m, p)	Community 4 (m, p)
Unipartite		(13, 1,208)	(9, 1,479)	(11, 1,187)	(1, 61)
Community 1 *(m, p)*	(3, 296)	(2, 134)	(0, 114)	(0, 5)	(1, 43)
Community 2 *(m, p)*	(11, 1,208)	(0, 44)	(2, 269)	*(9, 895)*	(0, 0)
Community 3 *(m, p)*	(2, 327)	(1, 137)	(1, 134)	(0, 52)	(0, 4)
Community 4 *(m, p)*	(8, 1,149)	(1, 51)	(5, 895)	(2, 193)	(0, 10)
Community 5 *(m, p)*	(10, 995)	*(9, 842)*	(1,67)	(0, 42)	(0, 4)

fully vaccinated, *partially vaccinated,* and *non-vaccinated* typifies this issue. For the mRNA vaccines, Pfizer and Moderna, does this mean 2 shots? And do boosters create a 3-shot vaccination program? Does this include a two-week post-shot period for the immune response to develop? For non-mRNA vaccines, *e.g.*, Johnson and Johnson, is it a 1-shot course or does it require a booster? And with a booster or multiple-shot protocols, how should the mixing of vaccines be considered as well as the length of time between shots? This complexity is further compounded when these events concern subjects in a frail state, even if temporary, as happens to women during pregnancy, *i.e.*, during the pre-conception, gestation, or post-partum periods, as to their impact on the baby or the mother. While the application of simple classifications facilitates statistical significance to be evaluated for such groupings, as Rose stated,[18] "…ideological reductionism manifests…confusion of statistical artifact with biological reality". Traditionally, there are two considerations in the use of statistical modeling to derive subgroup conclusions from the data. One assumes that the data are generated by a given stochastic data model. The other uses algorithmic models and treats the data dependency as being unknown.

We have used *community detection* to emphasize the difference between model-driven analysis and data-driven analysis. In the former, data is evaluated as to how well it may fit an existing model or how to refine the model to fit the data. In the latter, we are using an objective analysis of the data to identify what models may exist within the data.

Because of our focus on driving towards causality versus correlation and our intention to establish ease for clinical utility, we further incorporate clinical processes and pathways into the analysis. We distinguish our application of community detection from conventional unsupervised learning methods as we are developing and implementing complex functional models, *e.g.*, following clinical practice, to help identify and better manage potential gaps and biases that may exist in the data.

We note a fundamental problem in all analyses that attempt to identify "subgroups" of patients, namely that missing data may present a challenge that big data alone cannot overcome. We have several studies underway that focus on identifying symptoms that indicate missing critical data, although they do not necessarily identify what specific data may be missing. In addition, parallel development of the comprehensive functional model enables the identification of data that may bias analytic results and their interpretation.

The integration of a functional model that includes the complexity of disease processes, clinical practice, and the complexity of a patient, *e.g.*, lifestyle, environment, and genomics, is kept coherent by a knowledge graph that supports ongoing evolution as new concepts and relationships are identified, as well as a data model for functional integration of data from any source. The community detection algorithms described in this work operate on this knowledge graph.

Data enables an objective evaluation of its contents or any subset that may be selected to evaluate specific hypotheses. For example, communities can be identified that are based on data collected from EHRs and be compared with communities based on claims data to highlight the difference in perspective of these two data sources and their ability to describe a patient and their disease. Such a comparison could provide significant value to both the clinical community for improving patient management and the payer community for improving reimbursement policies, with both efforts yielding better outcomes for the patient.

While the application of our community detection approach introduced in this study has been focused on healthcare and a specific disease, HFpEF, it should be apparent that it represents an expanded view of how to address complex systems both in medicine and in many other fields. Community detection's ability, with properly binned data, to discover and return the medical variables that define each community, rather than them being specified by an analyst, delivers a valuable *unsupervised* capability. The integration of community detection algorithms with a model of the true complexity of the problem space should be viewed as being potentially generalizable across many domains and not limited to medicine, although it can provide a unique opportunity to improve clinical care and patient outcomes across most diseases and conditions.

Strengths and limitations

The potential clinical meaning underpinning the found communities has not been validated. Based on evidence for existing both liver and kidney pathology in patients with chronic liver disease, we defined the "Liver & Kidney data model" instead of two separate ones, a "Liver data model" and a "Kidney data model". Moreover, the absence of specific exam measurements suggested the combining of medical variables in a single data model. Specific exams for the Liver (Gamma Glutamyl Transferase, Glucose, ALP Alkaline phosphatase) were not present in the original data; other specific values for Kidney (Bicarbonate, Chloride, Urate) were absent too. Lastly, potential time-points for the longitudinal analysis were excluded because of many missing data among the different periods of the trial. By running frequencies on laboratory values collected in different periods (Baseline, Week 2 and 8; Month 6, 10, 14, 18, 30, 42, 54, 66, 72) we found the highest frequencies at Baseline and Month 72. For this reason, longitudinal analysis was done considering two timepoints only, presenting weaknesses in the sustainability of the results.

Despite the limitations above, we found that community detection techniques applied in this study are well suited to analyze complex phenomena involving large amounts of information; we had a significant number of subjects with associated data, approximately 4,000, and with many (complex) relationships to disease characteristics about these subjects. The application of community detection methods has identified critical aspects derived from overconnected portions of the network; the assessment of the quality of the resulting aggregation involved consultation with clinicians and domain experts.

Future directions

We believe that the results of this study can be generalized for redefining the concept of phenotype to incorporate the patient's progression through disease. We are currently applying this approach to several other diseases/conditions: multiple sclerosis, triple-negative breast cancer, and prematurity/infant-maternal morbidity and mortality. We believe that the next stage for refinement of diagnosis and both patient and disease stratification will require such evolution from current practice both for research and improved patient management and outcome.

Conclusions

This study has examined the application of community detection approaches to clinical trial data using algorithms applied to the I-PRESERVE study data. The resulting objective identification of

specific subgroups, without the need to initially establish the number of subgroups or number and identity of medical variables to include in the analysis, is a significant strength of this approach. In this manner, the data drives the analysis, and in the common situation where sparse data might be involved, is not limited in its ability to establish initial hypotheses about potential subgroups. The ability to incorporate disparate data without pre-selection also appears to potentially support a more integrative approach to clinical data analysis that can be used to improve disease stratification. These characteristics have the potential to enhance both the analysis of completed clinical trials, but of even greater significance, contribute to better clinical trial design and rates of success. The goal of this study was to review current limitations in disease stratification based on the need to apply reductionist considerations to achieve statistical significance and the potential to establish a data-driven, graph analytics approach that can lead to new study designs and hypotheses. We are currently applying these approaches in several other conditions and exploring the use of these results to further stratify HFpEF in clinical studies and patient management.

Supporting information

Supplementary material for this article is available at https://doi.org/10.14218/ERHM.2021.00081.

Supplementary File 1. The variable abbreviation description.

Conflict of interest

ML is a technical advisor to QCI. ML has been an editorial board member of *Exploratory Research and Hypothesis Medicine* since February 2020. MW and SR were at the time of this work QCI employees. The authors have no other conflicts of interest to report.

Author contributions

ML envisioned the whole experiment and ensured the connection to clinical usefulness. SP, MF, LF, MS, and SM prepared the data and did the initial unipartite-mCD analysis. MW performed the later bipartite-mCD analysis. SR provided analytic coherence and relevance for quantum computers. All authors have made a significant contribution to this study and have approved the final manuscript.

Data sharing statement

The original data used in this work may be obtained from the I-

PRESERVE study authors. The curated data for community detection is available from the authors.

References

[1] Breiman L. Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author). Statist Sci 2001;16(3):199–231. doi:10.1214/ss/1009213726.

[2] Eshghi A, Haughton D, Legrand P, Skaletsky M, Woolford S. Identifying Groups: A Comparison of Methodologies. Journal of Data Science 2011;9(2):271–291. doi:10.6339/JDS.201104_09(2).0009.

[3] Newman ME, Girvan M. Finding and evaluating community structure in networks. Phys Rev E Stat Nonlin Soft Matter Phys 2004;69(2 Pt 2):026113. doi:10.1103/PhysRevE.69.026113, PMID:14995526.

[4] Karataş A, Şahin S. Application Areas of Community Detection: A Review. 2018 International Congress on Big Data, Deep Learning and Fighting Cyber Terrorism (IBIGDELFT) 2018:65–70. doi:10.1109/IBIGDELFT.2018.8625349.

[5] Papadopoulos S, Kompatsiaris Y, Vakali A, Spyridonos P. Community detection in social media. Data Mining and Knowledge Discovery 2012;24(3):515–554. doi:10.1007/s10618-011-0224-z.

[6] Urbich M, Globe G, Pantiri K, Heisen M, Bennison C, Wirtz HS, *et al*. A Systematic Review of Medical Costs Associated with Heart Failure in the USA (2014-2020). Pharmacoeconomics 2020;38(11):1219–1236. doi:10.1007/s40273-020-00952-0, PMID:32812149.

[7] Hunt SA, Abraham WT, Chin MH, Feldman AM, Francis GS, Ganiats TG, *et al*. 2009 Focused update incorporated into the ACC/AHA 2005 Guidelines for the Diagnosis and Management of Heart Failure in Adults A Report of the American College of Cardiology Foundation/American Heart Association Task Force on Practice Guidelines Developed in Collaboration with the International Society for Heart and Lung Transplantation. J Am Coll Cardiol 2009;53(15):e1–e90. doi:10.1016/j.jacc.2008.11.013, PMID:19358937.

[8] Flores-Le Roux JA, Comin J, Pedro-Botet J, Benaiges D, Puig-de Dou J, Chillarón JJ, *et al*. Seven-year mortality in heart failure patients with undiagnosed diabetes: an observational study. Cardiovasc Diabetol 2011;10:39. doi:10.1186/1475-2840-10-39, PMID:21569580.

[9] Shah SJ, Borlaug BA, Kitzman DW, McCulloch AD, Blaxall BC, Agarwal R, *et al*. Research Priorities for Heart Failure with Preserved Ejection Fraction: National Heart, Lung, and Blood Institute Working Group Summary. Circulation 2020;141(12):1001–1026. doi:10.1161/CIRCULATIONAHA.119.041886, PMID:32202936.

[10] McMurray JJ, Carson PE, Komajda M, McKelvie R, Zile MR, Ptaszynska A, *et al*. Heart failure with preserved ejection fraction: clinical characteristics of 4133 patients enrolled in the I-PRESERVE trial. Eur J Heart Fail 2008;10(2):149–156. doi:10.1016/j.ejheart.2007.12.010, PMID:18279770.

[11] Partovi S, Trischman T, Kang PS. Lessons learned from the PRESERVE trial. Br J Radiol 2018;91(1087):20180092. doi:10.1259/bjr.20180092, PMID:29770709.

[12] Massie BM, Carson PE, McMurray JJ, Komajda M, McKelvie R, Zile MR, *et al*. Irbesartan in patients with heart failure and preserved ejection fraction. N Engl J Med 2008;359(23):2456–2467. doi:10.1056/NEJMoa0805450, PMID:19001508.

[13] McDonagh TA, Metra M, Adamo M, Gardner RS, Baumbach A, Böhm M, *et al*. 2021 ESC Guidelines for the diagnosis and treatment of acute and chronic heart failure. Eur Heart J 2021;42(36):3599–3726. doi:10.1093/eurheartj/ehab368, PMID:34447992.

[14] Bozkurt B, Hershberger RE, Butler J, Grady KL, Heidenreich PA, Isler ML, *et al*. 2021 ACC/AHA Key Data Elements and Definitions for Heart Failure: A Report of the American College of Cardiology/American Heart Association Task Force on Clinical Data Standards (Writing Committee to Develop Clinical Data Standards for Heart Failure). J Am Coll Cardiol 2021;77(16):2053–2150. doi:10.1016/j.jacc.2020.11.012, PMID:33250265.

[15] Mahmood SS, Levy D, Vasan RS, Wang TJ. The Framingham Heart Study and the epidemiology of cardiovascular disease: a historical perspective. Lancet 2014;383(9921):999–1008. doi:10.1016/S0140-6736(13)61752-3, PMID:24084292.

[16] Spilker B. Quality and quantity of data. Drug News Perspect 1998; 11(10):605–610. doi:10.1358/dnp.1998.11.10.863661, PMID:15616 629.

[17] Van Regenmortel MH. Reductionism and complexity in molecular biology. Scientists now have the tools to unravel biological and overcome the limitations of reductionism. EMBO Rep 2004;5(11):1016–1020. doi:10.1038/sj.embor.7400284, PMID:15520799.

[18] Rose S. What is Wrong with Reductionist Explanations of Behaviour? In: Bock GR, Goode JA (eds). Novartis Foundation Symposium 213 - The Limits of Reductionism in Biology. Hoboken: John Wiley & Sons Inc. 2007. doi:10.1002/9780470515488.ch13.

[19] Cleland JG, Swedberg K, Follath F, Komajda M, Cohen-Solal A, Aguilar JC, *et al*. The EuroHeart Failure survey programme— a survey on the quality of care among patients with heart failure in Europe. Part 1: patient characteristics and diagnosis. Eur Heart J 2003;24(5):442–463. doi:10.1016/s0195-668x(02)00823-0, PMID:12633546.

[20] Vaduganathan M, Michel A, Hall K, Mulligan C, Nodari S, Shah SJ, *et al*. Spectrum of epidemiological and clinical findings in patients with heart failure with preserved ejection fraction stratified by study design: a systematic review. Eur J Heart Fail 2016;18(1):54–65. doi:10.1002/ejhf.442, PMID:26634799.

[21] Solomon SD, Rizkala AR, Lefkowitz MP, Shi VC, Gong J, Anavekar N, *et al*. Baseline Characteristics of Patients With Heart Failure and Preserved Ejection Fraction in the PARAGON-HF Trial. Circ Heart Fail 2018;11(7):e004962. doi:10.1161/CIRCHEARTFAILURE.118.004962, PMID:29980595.

[22] Anker SD, Butler J, Filippatos G, Ferreira JP, Bocchi E, Böhm M, *et al*. Empagliflozin in Heart Failure with a Preserved Ejection Fraction. N Engl J Med 2021;385(16):1451–1461. doi:10.1056/NEJMoa2107038, PMID:34449189.

[23] Sun Y, Wang N, Li X, Zhang Y, Yang J, Tse G, *et al*. Predictive value of H2 FPEF score in patients with heart failure with preserved ejection fraction. ESC Heart Fail 2021;8(2):1244–1252. doi:10.1002/ehf2.13187, PMID:33403825.

[24] National Academies of Sciences, Engineering, and Medicine. In: Grumbling E, Horowitz M (eds). Quantum Computing: Progress and Prospects. Washington, DC: The National Academies Press. 2019. doi:10.17226/25196.

[25] Negre CFA, Ushijima-Mwesigwa H, Mniszewski SM. Detecting multiple communities using quantum annealing on the D-Wave system. PLoS One 2020;15(2):e0227538. doi:10.1371/journal.pone.0227538, PMID:32053622.

[26] Hen I, Spedalieri FM. Quantum annealing for constrained optimization. Phys Rev Applied 2016;5(3):034007. doi:10.1103/PhysRevApplied.5.034007.

[27] Hadfield S, Wang Z, O'Gorman B, Rieffel EG, Venturelli D, Biswas R. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. Algorithms 2019;12(2):34. doi:10.3390/a12020034.

[28] Davis A, Gardner BB, Gardner MR. Deep South: A social anthropological study of caste and class. Columbia: University of South Carolina Press. 2009.

[29] Liu X, Murata T. Community Detection in Large-scale Bipartite Networks. Transactions of the Japanese Society for Artificial Intelligence

2010;25(1):16–24.

[30] Barber MJ. Modularity and community detection in bipartite networks. Phys Rev E Stat Nonlin Soft Matter Phys 2007;76(6 Pt 2):066102. doi:10.1103/PhysRevE.76.066102, PMID:18233893.

[31] Calderer G, Kuijjer ML. Community Detection in Large-Scale Bipartite Biological Networks. Front Genet 2021;12:649440. doi:10.3389/fgene.2021.649440, PMID:33968132.

[32] Ganji M, Seifi A, Alizadeh H, Bailey J, Stuckey PJ. Generalized Modularity for Community Detection. In: Appice A, Rodrigues P, Santos Costa V, Gama J, Jorge A, Soares C (eds). Machine Learning and Knowledge Discovery in Databases Cham: Springer. 2015:655–670. doi:10.1007/978-3-319-23525-7_40.

[33] Franchini M, Pieroni S, Fortunato L, Knezevic T, Liebman M, Molinaro S. Integrated information for integrated care in the general practice setting in Italy: using social network analysis to go beyond the diagnosis of frailty in the elderly. Clin Transl Med 2016;5(1):24. doi:10.1186/s40169-016-0105-6, PMID:27465019.

[34] Franchini M, Pieroni S, Fortunato L, Molinaro S, Liebman M. Polypharmacy among the elderly: analyzing the co-morbidity of hypertension and diabetes. Curr Pharm Des 2015;21(6):791–805. doi:10.2174/1381612820666141024150901, PMID:25341855.

[35] Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. J Stat Mech 2008;2008(10):P10008.

[36] Cherven K. Network graph analysis and visualization with Gephi. Birmingham: Packt Publishing Ltd. 2013.

[37] Chukwu U, Dridi R, Berwald J, Booth M, Dawson J, Le D, *et al*. Constrained-optimization Approach Delivers Superior Classical Performance for Graph Partitioning via Quantum-ready Method. 2020 IEEE High Performance Extreme Computing Conference (HPEC) 2020:1–6.

[38] Teplukhin A, Kendrick BK, Mniszewski SM, Tretiak S, Dub PA. Sampling electronic structure quadratic unconstrained binary optimization problems (QUBOs) with Ocean and Mukai solvers. PLoS One 2022;17 (2):e0263849. doi:10.1371/journal.pone.0263849, PMID:35148343.

[39] Slack A, Yeoman A, Wendon J. Renal Dysfunction in Chronic Liver Disease. In: Vincent JL (ed). Yearbook of Intensive Care and Emergency Medicine 2010 Berlin, Heidelberg: Springer. 2010. doi:10.1007/978-3-642-10286-8_33.

[40] Mckelvie RS, Komajda M, McMurray J, Zile M, Ptaszynska A, Donovan M, *et al*. Baseline plasma NT-proBNP and clinical characteristics: results from the irbesartan in heart failure with preserved ejection fraction trial. J Card Fail 2010;16(2):128–134. doi:10.1016/j.cardfail.2009.09.007, PMID:20142024.

[41] Murari A, Peluso E, Lungaroni M, Gaudio P, Vega J, Gelfusa M. Data driven theory for knowledge discovery in the exact sciences with applications to thermonuclear fusion. Sci Rep 2020;10(1):19858. doi:10.1038/s41598-020-76826-4, PMID:33199734.

[42] Jalusic KO, Ellenberger D, Rommer P, Stahmann A, Zettl U, Berger K. Effect of applying inclusion and exclusion criteria of phase III clinical trials to multiple sclerosis patients in routine clinical care. Mult Scler 2021; 27(12):1852–1863. doi:10.1177/1352458520985118, PMID:33467978.

[43] Ayaz-Shah AA, Hussain S, Knight SR. Do clinical trials reflect reality? A systematic review of inclusion/exclusion criteria in trials of renal transplant immunosuppression. Transpl Int 2018;31(4):353–360. doi:10.1111/tri.13109, PMID:29274240.