RESEARCH

Stroke recovery phenotyping through network trajectory approaches and graph neural networks

Sanjukta Krishnagopal^{1*}, Keith Lohse² and Robynne Braun^{3,4}

*Correspondence:

s.krishnagopal@ucl.ac.uk

¹Gatsby Computational

Neuroscience Unit, University

College London, W1T 4JG,

London, UK

Full list of author information is
available at the end of the article

Abstract

Stroke is a leading cause of neurological injury characterized by impairments in multiple neurological domains including cognition, language, sensory and motor functions. Clinical recovery in these domains is tracked using a wide range of measures that may be continuous, ordinal, interval or categorical in nature, which presents challenges for standard multivariate regression approaches. This has hindered stroke researchers' ability to achieve an integrated picture of the complex time-evolving interactions amongst symptoms. Here we use tools from network science and machine learning that are particularly well-suited to extracting underlying patterns in such data, and may assist in prediction of recovery patterns. To demonstrate the utility of this approach, we analyzed data from the NINDS tPA trial using the Trajectory Profile Clustering (TPC) method to identify distinct stroke recovery patterns for 11 different neurological domains at 5 discrete time points. Our analysis identified 3 distinct stroke trajectory profiles that align with clinically relevant stroke syndromes, characterized both by distinct clusters of symptoms, as well as differing degrees of symptom severity. We then validated our approach using graph neural networks to determine how well our model performed predictively for stratifying patients into these trajectory profiles at early vs. later time points post-stroke. We demonstrate that trajectory profile clustering is an effective method for identifying clinically relevant recovery subtypes in multidimensional longitudinal datasets, and for early prediction of symptom progression subtypes in individual patients. This paper is the first work introducing network trajectory approaches for stroke recovery phenotyping, and is aimed at enhancing the translation of such novel computational approaches for practical clinical application.

Keywords: Stroke recovery; Disease subtyping; Network science; Network medicine; Graph Neural Networks

1 Introduction

Stroke Recovery is Dynamic and Involves Multiple Neurological Domains.

The process of neurological recovery after brain injuries such as stroke entails complex interactions among multiple variables that change dynamically over time [1, 2]. It is well known that the degree of recovery after stroke varies widely between individuals [3, 4, 5, 6], where each patient's recovery pattern uniquely reflects the combined influence of their lesion size and location [7], baseline health status, time to initial treatment [8], and response to medical treatment or rehabilitation, among many other intrinsic and extrinsic factors. Recovery trajectories furthermore vary depending on the specific neurological domain(s) affected (i.e. for motor, language, or sensory impairments), [9, 10] and each of these symptoms may show varying responsiveness to treatment. For example, language problems (aphasia), right-sided motor symptoms, and spatial perceptual problems (hemineglect) are reportedly less responsive than other symptoms to treatment with tissue plasminogen activator (tPA) [11]. Stroke recovery is therefore notoriously heterogeneous in terms of the type and severity of residual symptoms, as well as the timecourse of progression and/or resolution of those symptoms [12].

An important goal for stroke research is to reduce the 'noise' arising from this inherent heterogeneity by stratifying patients who are likely to have similar symptom trajectories. The heterogeneity of symptoms and time-varying recovery patterns inherent to stroke make it an area especially well-suited area for data-driven approaches. The increasing availability of large scale stroke datasets has led to a recent explosion in the use of data-science methods for stroke research [13]. For example, machine learning analyses of stroke clinical data [14, 12] have been used to characterize symptom clusters [15], predict outcomes [3], and define composite measures of recovery [16].

Limitations of Conventional Regression and Machine Learning Approaches versus Network Science

While Machine Learning (ML) approaches have been successful in a variety of analytical tasks, they often present challenges for interpretation and subsequent application. By contrast, network science tools are explicit in their modeling, making them more useful for studying medical data where clinical interpretation is

paramount. Additionally, typical ML approaches focus on prediction, while taking the outcome itself at face value; in contrast, network approaches attempt to improve how the outcome itself is captured. While ML tools may not be as readily interpretable as network approaches for various types of analyses (such as understanding the interactions that underlie disease recovery patterns) ML still presents several desirable properties, particularly in terms of data-driven predictive ability, and may therefore be useful for prognostication of patient recovery patterns.

Conventionally, statistical tools such as mixed-effects regression are used for modeling longitudinal data in disciplines where repeated measures designs are particularly relevant, such as education, motor learning, and psychology ([17, 18, 19]). Mixed-effects models have tremendous flexibility in their ability to accommodate different types of study designs, especially for handling repeated measures and the clustering of observations within groups. Such models are thus increasingly used in fields like neurorehabilitation where serial measures of recovery constitute a central focus [18, 20, 2]. Sophisticated though mixed-effects models may be as statistical tools, they face several challenges when applied to problems that commonly arise in recovery research, such as multivariate outcomes, ordinal outcomes, and violations of modeling assumptions. Many of these challenges can be readily addressed by using network approaches to data analysis.

Network Science and Trajectory Profile Clustering for Stroke Research

Insights into the complex patterns of symptom evolution can be gained through the computational power of networks analysis. The field of network medicine [21] studies disease manifestation and progression as a function of multiple interacting disease variables, which may be of similar or different types. Networks approaches also produce intuitive data visualizations that can facilitate interaction between clinicians and data scientists to yield novel insights on disease. However, with few exceptions, most network medicine studies have focused on biomolecular data [22, 23] rather than characterizing patients' patterns of symptom progression over time.

Recently, Krishnagopal et al. [24] introduced a network-based approach called Trajectory Profile Clustering (TPC) that groups patients based on similar patterns of symptom evolution. The intuitiveness and ability of TPC to integrate variables on multiple different scales make it especially useful for studying disease severity, progression, and recovery. Multi-layer [25] types of trajectory clustering have also

shown success in clinically-validated disease trajectory prediction in Parkinson's disease. We argue that TPC offers unique advantages for stroke recovery research based on its ability to simultaneously (i) identify the dominant variables that differentiate stroke recovery subtypes, (ii) account for temporal disease progression patterns, and (iii) delineate distinct symptom groupings. This paper is the first work introducing network trajectory approached for stroke recovery phenotyping, and is aimed at enhancing the translation of such novel computational approaches for practical clinical application.

When analyzing recovery trajectories with TPC, an obvious question that arises is at what stage do patients begin to stratify into distinct trajectory clusters (i.e. when do they begin to show symptom patterns unique to their recovery subtype)? The timing of medical treatments might be one important influence on the time-course of recovery subtype stratification. For example in stroke, stratification might be expected to occur based on when patients receive treatments such as tPA or clot retrieval. Naively, this may appear to be a problem of simply measuring the differences between trajectory subtypes at each timepoint. However, since treatment efficacy for the same individual at different timepoints is not unrelated, more sophisticated tools are required to extract the timescale of separation. We can investigate these questions through graphical tools such as Graph Neural Networks in machine learning.

Graph Neural Networks for the Study of Neurological Disorders.

The field of machine learning has been revolutionized by recent advances in deep neural networks, especially convolutional neural networks (CNNs) [26]. Conventionally, CNNs use local connections, shared weights and multiple layers to extract representations of data. However, CNNs work in a Euclidian domain and are best suited for use with images. By contrast, other deep learning methods can operate on a graph domain (i.e. graph neural networks or GNNs) [27]. Convolutional variants of graph neural networks provide a framework for transferring deep learning operators into a non-Euclidean (graphical) domain, and have been successful in a variety of tasks such as graph classification, node identification, link prediction in protein interactions, knowledge graphs, and social network analysis, among others. Of particular interest here, they have been successfully applied to study neurologi-

cal disorders including Alzheimer’s disease and autism [28, 29], and could similarly have utility in the study of stroke.

2 Stroke Dataset Analyzed

To demonstrate the utility of using TPC and CNN for stroke recovery research, we analyzed cases from the well-characterized NINDS tPA trial data set [30]. This study was a randomized, double-blinded, placebo-controlled trial that compared the effects of intravenous tPA (a thrombolytic agent used in ischemic stroke to dissolve blood clots) versus placebo treatment in 624 patients. The data set captures neurologic deficits on the NIH Stroke Scale (NIHSS) [31], which is the most widely used measurement scale for stroke neurologic deficits, and has clinimetric properties that are well-defined [32, 33]. Each item is scored on a scale (from 0-3 or 0-5), with higher values indicating greater stroke severity. The NINDS tPA data captures NIHSS scores across 5 time points: at hospital admission, at 2 hours, 24 hours, 7-10 days, and 3 months post-stroke. Here, we examined symptom progression in 11 neurologic domains as assessed by 15 individual item subscores on the NIHSS. We excluded a total of 135 cases who had imputed data at any time point (134) and/or had died (118). We excluded these cases because the imputation approach that had been used could distort patterns of change in scores for individual patients (i.e. missing values were imputed as the worst score for each NIHSS item). After exclusions, there were 489 remaining cases for analysis.

Table 1 NINDS tPA Trial Data, Variable Names and Symptom Descriptions

Variable Name	Symptom Description
ATAXIA	Coordination
CONSCIO	Level of consciousness
DYSAR	Speech (slurring)
EXTIN	Spatial perception
GAZE	Eye movements
LANG	Language
LOCCOM	Command following
LOCQU	Question answering
MOTORLA	Left arm strength
MOTORLL	Left leg strength
MOTORRA	Right arm strength
MOTORRL	Right leg strength
PALSY	Facial weakness
SENSORY	Skin sensation
VISUAL	Visual fields

3 Methods

3.1 Mixed-Effects Regression Model

The NINDS dataset contains repeated measures at 5 time points across 11 different neurological domains (as measured by 15 different NIHSS assessment items), resulting in 75 observations per patient. To understand how changes in different symptoms related to each other over time, we extracted the random-effects from the model to compute slopes and intercepts for each individual overall, and for each individual in each domain. Using these coefficients we could test two different hypotheses to better understand behavioral phenotypes: (1) how slopes related to intercepts both overall and within specific domains and (2) how slopes (rates of change) related to each other across domains (i.e. which symptoms improve together). Given that the NIHSS items have different maxima (all have a minimum of 0), we choose to standardize NIHSS scores to be a fixed percentage of the maximum in each domain. Fig. 1 shows these standardized data for each NIHSS assessment item as a function of time. Inspection of the raw data revealed substantial right skew in this dependent variable, so we log-transformed the standardized values ($\log_{10}(x_{hij}/\max(x_h) + 1)$),

to achieve a more normal distribution of scores. (Note, h indexes the scale, i the individual participant, and j the specific time-points).

To account for the statistical dependence among these repeated measures, we include random-intercepts and random-slopes for the effect of ‘Time’ at both the level of patient, and at the level of patient:domain. The technical details are outlined here. Conceptually, effects at the patient level reflect the average change for a patient over time across domains, represented by a unique slope and intercept for each individual, $(1 + Time|patient)$. Effects at the domain-level reflect the change over time in each domain for a particular patient, which can be captured by a unique slope and intercept in each domain for each individual, $(1 + Time|(patient : domain))$. Models also includes fixed-effects of Time, Group (tPA versus placebo), Domain (the 15 assessment items of the NIHSS), and their interactions. We test the normality of the distribution of residuals and the random-effects at each level using the Kolmogorov-Smirnov test. Statistical significance for these effects is based on Type III Wald χ^2 -tests for the change in deviance ($\alpha = 0.05$). These χ^2 -tests and p-values thus reflect null-hypothesis significance tests controlling for all of the other variables in the model.

3.2 Trajectory Profile Clustering

The Trajectory Profile Clustering algorithm [24] is designed to group together patients based on the similarities of their disease trajectories. In essence, it uses graphical tools to generate trajectory profiles for each individual that track their evolution of symptoms across time, then clusters them into communities of similarly behaving individuals that define a recovery subtype. The algorithm proceeds as follows:

- 1 *Model using bipartite networks:* At time point t we construct a bipartite graph modeling connections between N individuals and V disease variables/ symptoms. The connections between the individuals and symptoms are encoded through an adjacency matrix A_t of size $N \times V$. For M time points, we can represent the set of these bipartite graphs as an $N \times V \times M$ by stacking the A_t across time points to generate a tensor X where $X_{i vt}$ gives the value of individual i ’s disease symptom v at time t .
- 2 *Threshold and binarize to obtain trajectory profile:* We threshold each symptom to set values less than a fixed fraction κ of the maximum score for

the symptom to zero,. For example, if symptom ν takes score values in $(0,1,2,3,4,5)$, if $\kappa = 0.5$, we binarize scores such that scores below $5 \times 0.5 = 2.5$ are set to 0, and above are set to 1. We call this thresholded matrix the trajectory profile matrix, T^i for patient i , that contains a representation of how the set of symptoms that a patient is *severely affected* by varies with time. The matrix entries of T^i are calculated as follows:

$$T_{\nu t}^i = 0 \text{ if } X_{i\nu t} \leq \max(\nu) \cdot \kappa \quad (1)$$

$$= 1 \text{ if } X_{i\nu t} > \max(\nu) \cdot \kappa \quad (2)$$

Since the range of values for each symptom represents the entire scale of severity, this thresholding ensures that patients are only considered ‘connected’ to symptoms that they severely express.

- 3 *Create a patient-patient network based on trajectory similarity:* We create a patient-patient network P of all patients. The nodes of this network denote patients, and the strength of a link between patient i and patient j captures the similarity of their trajectory profiles. P has an adjacency matrix given by:

$$P_{ij} = \sum_{v,t} (T_{vt}^i \equiv T_{vt}^j). \quad (3)$$

In other words, P_{ij} gives the number of matrix entries for which trajectory profile T^i has the same value as T^j . This formulation implies that symptoms are *equally weighted*. While the approach is amenable to non-uniform weighting, there is little clinical consensus on the relative importance of different symptoms. Hence, in the interest of not introducing external bias, we choose uniform weighting, adopting an agnostic approach that assumes all symptoms/indicators are equally important. Other applications may require unequal weighting for symptoms and different time points, in which case one may calculate the patient-patient matrix as follows: $P_{ij} = \sum_{v,t} w_{vt} (T_{vt}^i \equiv T_{vt}^j)$ where w_{vt} is the weight of symptom v at time t .

- 4 *Cluster the network to identify subtypes:* We then perform Louvain community detection [34] to maximize the Newman-Girvan modularity function [35] on the network defined by the adjacency matrix P . Such community detec-

tion allows us to identify ‘communities’ of patients, where individuals within a community have a relatively more similar stroke recovery profiles than patients between communities. As is common in network community detection approaches, the number of communities is not set a priori, but rather chosen so that the modularity is maximized. This process allows us to cluster trajectory profiles, and hence patients, into subtypes which have high intra-subtype similarity. The subtypes are denoted by $C^1, C^2, \dots, C^L$ where each C^l is a collection of trajectory profiles of the patients in that subtype, and L is the total number of subtypes.

- 5 *Construct aggregate profiles to characterize each subtype:* We average the trajectory profiles of all patients in each community C^l to obtain the ‘community/subtype profile’ S^l . The subtype profile is indicative of the symptom features that describe the subtype. More specifically, it is the normalized average of the trajectory profiles of all the patients in that subtype, i.e., S^l is a $V \times M$ matrix with elements defined by

$$S_{vt}^l = \frac{\sum_{i \in C^l} T_{vt}^i}{N_l} \quad (4)$$

where N_l is the total number of individuals in subtype C^l .

3.3 Graph Neural Network

Graph neural networks (GNNs) [27] are a machine learning approach that captures the relationships represented in graphs through message passing between the nodes of those graphs. GNNs take a graph as input and pass them several layers of nodes, artificial ‘neurons’. Here we use graph neural networks to identify the timepoints that are most relevant in determining stroke recovery subtypes. Specifically, we train a Graph Neural Network on symptom-symptom graphs generated at each timepoint, and test the accuracy of a GNN in its ability to classify an individual into the correct recovery subtype using data from a single timepoint. A higher accuracy at a given timepoint implies that the recovery subtypes attributed to the patients are strongly correlated with the symptom profiles at that timepoint.

We generate a symptom-symptom interaction graph G_t at each timepoint t where the nodes represent the disease symptoms. This graph is undirected (i.e., if node x is

connected to node y , vice-versa is also true, i.e., the adjacency matrix of this graph is symmetric). The graph is generated as follows. First we generate a symptom-patient binary interaction network for a given timepoint as in step (3) of the previous section. We then project it to symptom space to obtain the symptom interaction profile of each patient at a timepoint. The corresponding adjacency matrix (of size $V \times V$) for the graph for patient i is given by

$$G_{it} = T_{ivt}^T \times T_{ivt} \quad (5)$$

Lastly, We repeat this for each individual such that there exist N symptom-symptom interaction graphs at each timepoint.

We then separate the individual cases into a training data set (70% of total individuals) and a test data set (30% of total individuals) used to validate our approach. A convolutional graph neural network is trained on a graph classifying task for each time point, with labels provided by the subtypes/communities of that individual. The stratification of individuals into their recovery subtypes at each timepoint is then measured by testing the accuracy of the GNN on the test data for each timepoint.

The graph neural network takes as input symptom-symptom networks where we consider the 15 NIHSS assessment items as the nodes. The network consists of an input layer, a single hidden layer, and an output layer. The hidden layer comprises of 64 artificial neurons. The input is processed through two graph convolutional layers with ReLU nonlinearities. We then calculate the graph representation by averaging all the neuron representations in the output layer, which contains an equal number of neurons to the number of subtypes. The output is passed through a softmax classifier that yields the probability of the graph belonging to a particular category/subtype. We use cross-entropy loss and an optimizer (ADAM) for adaptive moment estimation.

4 Results

4.1 Key Patterns in Slopes from the Linear Mixed-Effect Models

First, we present results of the linear mixed-effects-model, a more conventional statistical approach used to model medical data. Modeling each domain of the NIHSS using the mixed-effects model demonstrated several key-patterns in the data. First,

the level of impairment at the beginning of the trial (i.e., at the intercept) differed by domain. We measure the chi-squared value to test if the differences between domains were statistically significant, and found a main-effect of domain $\chi^2(14) > 999, p < 0.001$, a main-effect of Time, $\chi^2(1) = 677.48, p < 0.001$, and there was a significant Time x Domain interaction, $\chi^2(14) = 609.42, p < 0.001$, showing that impairments in different domains resolved at different rates over time, as shown in Fig. 1. On average across time and domains, there was also a statistically significant main-effect of treatment vs placebo, $\chi^2(1) = 12.79, p < 0.001$. There was some evidence for a treatment Group x Domain x Time interaction, $\chi^2(14) = 20.83, p = 0.106$ which, although not statistically significant, is conceptually consistent with prior analyses of this data [11] showing that the benefits of tPA differ depending on the domain of impairment and the time-post stroke.

Inspection of the random effects of this model also provides some insights relevant to the current goal of creating behavioral phenotypes. As shown in Fig. 2(A), there was a negative association between slopes and intercepts, on average across the NIHSS assessment items, such that *subjects who had higher levels of initial impairment had more negative slopes*. This negative association is (at least in part) driven by a floor effect in the NIH stroke scale, where zero values indicate an assessment of no impairment. This pattern was generally consistent across the different domains of the NIHSS.

More importantly, as shown in Fig. 2(B), there were several unique correlations among slopes at the patient:domain level of the model that are of interest: First, as would be expected given the common occurrence of post-stroke hemiplegia (weakening on one side of the body), some of the strongest correlations were for the ipsilateral arm and leg. The right arm and leg showed a similar time-course of change in impairment ($r = 0.74$), as did the left arm and leg ($r = 0.79$). Second, there are also patterns in the correlation matrix consistent with lateralization of function as affected by unilateral stroke. For instance, the NIHSS item for extinction was positively associated with left arm/leg deficits (i.e. right hemisphere damage), and negatively associated with right arm/leg deficits (i.e. left hemisphere damage). Gaze was also positively associated with this set of symptoms, albeit to a much smaller degree, possibly reflecting the fact that gaze deviation is typically more pronounced in patients with hemineglect [36]. Lastly, language was positively associated with

the right arm/leg (left hemisphere) and negatively associated with the left arm/leg (right hemisphere).

The three points outlined above highlight critical insights that can be provided by conventional approaches. However, this linear mixed-effects models is unable to resolve a finer level of detail for the time-evolving interactions that may underlie stroke recovery subtypes. Additionally, limitations in interpreting these data are the non-normal distributions of both the residuals and the random-effects from the model. For instance, a Kolmogorov-Smirnov test of the studentized residuals showed a strong departure from normality, $p < 2.2e-10$. Thus, although the marginal effects are informative, this statistical approach is clearly not optimal for understanding how patients' symptoms are changing over time with these ordinal data where strong ceiling/floor effects are present. These limitations are overcome in the network-based TPC method, which is particularly well suited for longitudinal data of any type (e.g. normal, non-normal categorical, ordinal). The agreement between the two approaches, as detailed below, serves to further validate TPC as an effective, flexible and intuitive approach.

4.2 Stroke Recovery Subtypes identified by TPC

Maximizing modularity on the patient-patient trajectory-similarity network gives us three distinct recovery subtypes. It is worth mentioning that the number of

subtypes are not predetermined, but are optimally chosen such that the modularity is maximized, i.e., the subtypes are optimally separated.

Fig.3 shows the clinical profiles of each subtype. The darkness of the shade of grey for each symptom over time denotes the fraction of patients who had a value above threshold for that symptom. To reiterate, the thresholding ensures that patients are only considered affected with symptoms for which they display relatively high severity, defined to be above the population median. Our analysis identifies 3 distinct stroke trajectory profiles that align with clinically relevant stroke syndromes, characterized both by distinct clusters of symptoms, as well as differing degrees of symptom severity over time. Several key features of the identified subtypes warrant comment. First, our TPC approach identifies a 'mildly affected' group that was the least symptomatic of the three subtypes both in terms of baseline severity and

3-month residual symptoms. As a group, this subtype showed a mixture of features that are not clearly lateralizing. In addition, two severely affected subtypes are readily identified that correspond to left and right hemisphere syndromes: We find a ‘left motor’ subtype, showing severely impaired left arm and leg strength together with hemineglect, but with essentially no right-sided motor symptoms (red boxes, Fig.3), and a ‘right motor’ subtype, showing severely impaired right arm and right leg strength together with aphasia (blue boxes, Fig.3). Additionally, spatial perception scale items are most affected in the ‘left motor’ group (corresponding clinically to a right hemisphere syndrome with hemispatial neglect). Conversely, the language and question-answering items are most affected in the ‘right motor’ group (corresponding to a left hemisphere syndrome with primarily expressive aphasia). These findings are in alignment with prior factor analysis on the clinimetric properties of the NIHSS [37] and Principal Component Analysis (PCA) to define common behavioral clusters [38], as well as results from our mixed-effects model in section 4.1. The fact that our results capture clinically relevant subtypes and corroborate these prior findings supports the content validity of this analytic approach. Additionally, our clustering approaches reveal subtype structure at a finer scale (both in terms of symptoms as well as longitudinal symptom evolution) than can be achieved with PCA, and results that are clinically consistent.

It is notable that all three identified trajectory subtypes included both tPA- and placebo-treated patients, suggesting that treatment effects were less defining characteristics of patient recovery profiles than were initial severity and stroke laterality. TPC also provides interesting insights into patterns of symptom prevalence over time across subtypes. Spatial perception deficits (hemineglect) are present in both the left-motor and right-motor subtypes, but tend to be milder and have better resolution in left hemisphere strokes. This observation reinforces the importance of targeted screening during rehabilitation for hemineglect symptoms in both left- and right-hemisphere stroke, since persistent milder symptoms that could be amenable to treatment might otherwise be overlooked. Visual deficits are also present in both left- and right-hemisphere strokes as would be expected, but contrary to conventional understanding that visual deficits resolve less well than hemineglect, the overall prevalence of persistent visual symptoms at 3 months is lower than for hemineglect.

4.2.1 Effects of tPA Treatment and Time

A natural extension of our TPC subtyping is to study the timecourse of stratification into these subtypes. One might wonder whether subtype (and consequently the expected recovery profile of a patient) is largely driven by baseline symptom severity or by early stroke treatments. Machine learning is particularly well-posed to answer such questions. Since we operate in the graph domain, we use graph neural networks. We first extract trajectory profiles independently on patients who had received tPA

within 3 hours of stroke onset compared to patients who had received a placebo. In Fig. 4 we see that the identified subtypes retain the same symptom clusters identified in Figure Fig.3(B), but that overall symptom severity is lower in the tPA-treated population Fig. 4(A), particularly in symptoms that are dominant identifiers of the group. For instance, in the left-motor group, assessment items for gaze, left arm and leg strength, and sensation showed higher recovery in the tPA-treated group (Fig. 4(A)) versus placebo group (Fig. 4(B)). Similarly, in the right-motor group, assessment items for right arm and leg strength, and command-following (SLOC-COM) showed higher recovery in the tPA-treated group (Fig. 4(A)) compared to the placebo group (Fig. 4(B)). As expected, the effect of tPA is less obvious in the minimally impaired group. We explore in Fig. 5 the accuracy of a neural network in predicting the subtype of an individual given data at a discrete timepoint. We generated a symptom-interaction network for each individual at each timepoint, and trained a convolutional-GNN to learn properties of interaction with neighbors. The convolutions are used for averaging over the neighborhood. If the learned properties for that timepoint are separated according to subtypes, then data at that timepoint is considered a good predictor of the subtype. In the training stage, we assume that the subtype is known to the neural network, which attempts to learn correlations between symptom-interaction patterns and the subtype. We then test to identify if the features learned by the neural network are consistent with the actual subtypes of the test patients.

Fig. 5 shows that there is a difference in predictive accuracy at baseline for the tPA vs. Placebo groups. The baseline timepoint is a more accurate predictor of subtype for patients who received tPA. This finding may seem unexpected if one posits that tPA ‘rescues’ an otherwise poor prognosis with severe baseline symptoms predicting poor outcomes. However the predictive accuracy for the tPA group rapidly increases during the first 2 hours, and given the expected timecourse for therapeutic effects of tPA, this finding provides additional validation of our approach. For the Placebo group, predictive accuracy grows at comparable rates (comparable slopes) up to the 24 hour mark, showing peak predictive accuracy at the 7-10 day mark, with standard error on the order 10^{-2} . The tPA group showed a further uptick in predictive accuracy by the 3 month mark, suggesting that treatment continues to exert an effect on recovery subtype stratification even in the later stages of post-stroke. One

might speculate that this is the result of tPA treatment salvaging a greater ‘reserve’ of neural tissue for later rehabilitation therapies to act upon. Our report on the rapid increase in predictive accuracy from 2-24 hours post-stroke furthermore aligns with recent work by Heitsch *et al.* [39] who reported on the early change in NIHSS scores between 6 – 24 hours as a dynamic phenotype associated with long-term outcomes.

5 Conclusion and Future Work

In this work we introduce a network-based, data-driven method for stroke recovery analysis. First, we analyze the NINDS tPA stroke dataset using conventional quantitative medicine methods including a linear mixed-effects regression model, examining the effects of time, group (tPA vs. Placebo), neurological domain, and their interactions. Further, to identify stroke recovery subtypes and examine their characteristics at a finer resolution, we use the Trajectory Profile Clustering method which accounts not only for symptom severity at different timepoints, but also symptom interactions and their temporal evolution. Of note, although the analytical approach is clinically agnostic, we identify subtypes that are clinically relevant. In particular, we identify a mildly-affected recovery subtype comprising a larger proportion of patients who received tPA. Additionally, we observed that the two other recovery subtypes stratify as left- versus right-sided hemiplegia. Additionally, we identified that left motor deficits are strongly correlated with deficits in gaze and extinction, whereas right motor deficits correlated with deficits in language. These results again are biologically relevant, and are further validated by convergent findings in the mixed-effects regression model. Lastly, we use graph neural networks to study how much of the stratification into subtypes is identifiable at different time points, and found that stroke recovery trajectories were largely defined within the first 24 hours, consistent with the expected pharmacodynamics of tPA treatment delivered in the first 3 hours after stroke.

This paper is the first work introducing network trajectory approaches for stroke recovery phenotyping, and is aimed at enhancing the translation of such novel computational approaches for practical clinical application. This work presents a data-driven method that is widely applicable to heterogeneous neurological disorders such as stroke, and bridges the fields of predictive medicine and network informatics. Because our approach is uniquely adapted to accommodate input variables on multiple

scales, future applications could include the integration of other types of data that may contribute to the heterogeneity of recovery, such as data on patient genotypes.

Availability of Data

The NINDS dataset was first released in [40] and is publicly available. The code is available on [github/chimeraki/Stroke-Analysis](#)

Competing interests

The authors declare that they have no competing interests.

Author's contributions

SK developed and implemented the TPC algorithm. KL implemented and studied the mixed-effects model on the data. RB conceived, facilitated and guided the project. SK and KL analyzed the computational results and KL and RB interpreted the results in a medical context. All authors were involved in writing the manuscript.

Acknowledgements

We would like to acknowledge Michelle Girvan, Adam de Havenon, Michael M. Binkley, Laura Heitsch, and Alen Delic for helpful conversations and perspectives. We would also like to thank Steven C. Cramer and Arne G. Lindgren for their diligent critical review and comments on the manuscript.

Author details

¹Gatsby Computational Neuroscience Unit, University College London, W1T 4JG, London, UK. ²Physical Therapy and Neurology, Washington University School of Medicine, 4444 Forest Park Ave., Suite 1101, 63108-2212 St. Louis, MO, USA. ³Department of Neurology, Baltimore, Bressler Research Building, University of Maryland, 655 W. Baltimore Street, 12th Floor, 21201 Baltimore, MD, USA. ⁴GPAS Collaboration, Phenotyping Core.

References

- Wang, C., Winstein, C., D'Argenio, D.Z., Schweighofer, N.: The efficiency, efficacy, and retention of task practice in chronic stroke. *Neurorehabilitation and neural repair* **34**(10), 881–890 (2020)
- Lohse, K., Bland, M.D., Lang, C.E.: Quantifying change during outpatient stroke rehabilitation: a retrospective regression analysis. *Archives of physical medicine and rehabilitation* **97**(9), 1423–1430 (2016)
- Stinear, C.M., Barber, P.A., Petoe, M., Anwar, S., Byblow, W.D.: The PREP algorithm predicts potential for upper limb recovery after stroke. *Brain: A Journal of Neurology* **135**(Pt 8), 2527–2535 (2012). doi:10.1093/brain/aws146
- Wahl, A.-S., Schwab, M.E.: Finding an optimal rehabilitation paradigm after stroke: enhancing fiber growth and training of the brain at the right moment. *Frontiers in Human Neuroscience* **8**, 381 (2014). doi:10.3389/fnhum.2014.00381
- Simpkins, A.N., Janowski, M., Oz, H.S., Roberts, J., Bix, G., Doré, S., Stowe, A.M.: Biomarker Application for Precision Medicine in Stroke. *Translational Stroke Research* **11**(4), 615–627 (2020). doi:10.1007/s12975-019-00762-3. Accessed 2021-07-05
- Braun, R.G., Heitsch, L., Cole, J.W., Lindgren, A.G., de Havenon, A., Dude, J.A., Lohse, K.R., Cramer, S.C., Worrall, B.B., Core, P., *et al.*: Domain-specific outcomes for stroke clinical trials: What the modified rankin isn't ranking. *Neurology* **97**(8), 367–377 (2021)
- Shelton, F.D.N., Reding, M.J.: Effect of lesion location on upper limb motor recovery after stroke. *Stroke* **32**(1), 107–112 (2001)
- Dromerick, A.W., Edwardson, M.A., Edwards, D.F., Giannetti, M.L., Barth, J., Brady, K.P., Chan, E., Tan, M.T., Tamboli, I., Chia, R., *et al.*: Critical periods after stroke study: translating animal stroke recovery experiments into a clinical trial. *Frontiers in human neuroscience* **9**, 231 (2015)
- Cramer, S.C., Koroshetz, W.J., Finklestein, S.P.: The case for modality-specific outcome measures in clinical trials of stroke recovery-promoting agents. *Stroke* **38**(4), 1393–1395 (2007)
- Cramer, S.C., Wolf, S.L., Saver, J.L., Johnston, K.C., Mocco, J., Lansberg, M.G., Savitz, S.I., Liebeskind, D.S., Smith, W., Wintermark, M., Elm, J.J., Khatri, P., Broderick, J.P., Janis, S., NIH StrokeNet Recovery and

- Rehabilitation Group and the Acute Stroke Group*: The utility of domain-specific end points in acute stroke trials **52**(3), 1154–1161. doi:10.1161/STROKEAHA.120.031939
11. Felberg, R.A., Okon, N.J., El-Mitwalli, A., Burgin, W.S., Grotta, J.C., Alexandrov, A.V.: Early dramatic recovery during intravenous tissue plasminogen activator infusion: clinical pattern and outcome in acute middle cerebral artery stroke **33**(5), 1301–1307. doi:10.1161/01.str.0000015556.48283.74
12. van der Vliet, R., Selles, R.W., Andrinopoulou, E.-R., Nijland, R., Ribbers, G.M., Frens, M.A., Meskers, C., Kwakkel, G.: Predicting upper limb motor impairment recovery after stroke: a mixture model. *Annals of neurology* **87**(3), 383–393 (2020)
13. Wang, W., Kiik, M., Peek, N., Curcin, V., Marshall, I.J., Rudd, A.G., Wang, Y., Douiri, A., Wolfe, C.D., Bray, B.: A systematic review of machine learning models for predicting outcomes of stroke with structured data. *PloS one* **15**(6), 0234722 (2020)
14. Stinear, C.M., Byblow, W.D., Ackerley, S.J., Smith, M.-C., Borges, V.M., Barber, P.A.: Prep2: A biomarker-based algorithm for predicting upper limb function after stroke. *Annals of clinical and translational neurology* **4**(11), 811–820 (2017)
15. Sucharew, H., Khoury, J., Moomaw, C.J., Alwell, K., Kissela, B.M., Belagaje, S., Adeoye, O., Khatri, P., Woo, D., Flaherty, M.L., Ferioli, S., Heitsch, L., Broderick, J.P., Kleindorfer, D.: Profiles of the National Institutes of Health Stroke Scale items as a predictor of patient outcome. *Stroke* **44**(8), 2182–2187 (2013). doi:10.1161/STROKEAHA.113.001255
16. Hommel, M., Detante, O., Favre, I., Touzé, E., Jaillard, A.: How to Measure Recovery? Revisiting Concepts and Methods for Stroke Studies. *Translational Stroke Research* **7**(5), 388–394 (2016). doi:10.1007/s12975-016-0488-0
17. Singer, J.D., Willett, J.B., Willett, J.B., *et al.*: *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence*. Oxford university press, ??? (2003)
18. Garcia, T.P., Marder, K.: Statistical approaches to longitudinal data analysis in neurodegenerative diseases: Huntington’s disease as a model. *Current neurology and neuroscience reports* **17**(2), 14 (2017)
19. Lohse, K.R., Shen, J., Kozlowski, A.J.: Modeling longitudinal outcomes: A contrast of two methods. *Journal of motor learning and development* **8**(1), 145–165 (2020)
20. Kozlowski, A.J., Heinemann, A.W.: Using individual growth curve models to predict recovery and activities of daily living after spinal cord injury: an scirehab project study. *Archives of physical medicine and rehabilitation* **94**(4), 154–164 (2013)
21. Pawson, T., Linding, R.: Network medicine. *FEBS letters* **582**(8), 1266–1270 (2008)
22. Huang, S.: Back to the biology in systems biology: what can we learn from biomolecular networks? *Briefings in Functional Genomics* **2**(4), 279–297 (2004)
23. Barabási, A.-L., Gulbahce, N., Loscalzo, J.: Network medicine: a network-based approach to human disease. *Nature reviews genetics* **12**(1), 56–68 (2011)
24. Krishnagopal, S., Coelln, R.v., Shulman, L.M., Girvan, M.: Identifying and predicting parkinson’s disease subtypes through trajectory clustering via bipartite networks. *PloS one* **15**(6), 0233296 (2020)
25. Krishnagopal, S.: Multi-layer trajectory clustering: A network algorithm for disease subtyping. *Biomedical Physics & Engineering Express* **6**(6), 065003 (2020)
26. Albawi, S., Mohammed, T.A., Al-Zawi, S.: Understanding of a convolutional neural network. In: 2017 International Conference on Engineering and Technology (ICET), pp. 1–6 (2017). Ieee
27. Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. *IEEE transactions on neural networks* **20**(1), 61–80 (2008)
28. Song, T.-A., Chowdhury, S.R., Yang, F., Jacobs, H., El Fakhri, G., Li, Q., Johnson, K., Dutta, J.: Graph convolutional neural networks for alzheimer’s disease classification. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), pp. 414–417 (2019). IEEE
29. Parisot, S., Ktena, S.I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., Rueckert, D.: Disease prediction using graph convolutional networks: application to autism spectrum disorder and alzheimer’s disease. *Medical image analysis* **48**, 117–130 (2018)
30. National Institute of Neurological Disorders and Stroke rt-PA Stroke Study Group: Tissue plasminogen activator for acute ischemic stroke. *The New England Journal of Medicine* **333**(24), 1581–1587 (1995).

- doi:10.1056/NEJM199512143332401
31. Brott, T., Adams, H.P., Olinger, C.P., Marler, J.R., Barsan, W.G., Biller, J., Spilker, J., Holleran, R., Eberle, R., Hertzberg, V.: Measurements of acute cerebral infarction: a clinical examination scale. *Stroke* **20**(7), 864–870 (1989). doi:10.1161/01.str.20.7.864
 32. Millis, S.R., Straube, D., Iramaneerat, C., Smith, E.V., Lyden, P.: Measurement properties of the National Institutes of Health Stroke Scale for people with right- and left-hemisphere lesions: further analysis of the clomethiazole for acute stroke study-ischemic (class-I) trial. *Archives of Physical Medicine and Rehabilitation* **88**(3), 302–308 (2007). doi:10.1016/j.apmr.2006.12.027
 33. Goldstein, L.B., Samsa, G.P.: Reliability of the national institutes of health stroke scale: extension to non-neurologists in the context of a clinical trial. *Stroke* **28**(2), 307–310 (1997)
 34. Blondel, V.D., Guillaume, J.-L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* **2008**(10), 10008 (2008)
 35. Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proceedings of the national academy of sciences* **99**(12), 7821–7826 (2002)
 36. Deviation of eyes and head in acute cerebral stroke **6**. doi:10.1186/1471-2377-6-23
 37. Lyden, P., Lu, M., Jackson, C., Marler, J., Kothari, R., Brott, T., Zivin, J.: Underlying structure of the national institutes of health stroke scale: results of a factor analysis. *NINDS tPA stroke trial investigators* **30**(11), 2347–2354. doi:10.1161/01.str.30.11.2347
 38. Corbetta, M., Ramsey, L., Callejas, A., Baldassarre, A., Hacker, C.D., Siegel, J.S., Astafiev, S.V., Rengachary, J., Zinn, K., Lang, C.E., Connor, L.T., Fucetola, R., Strube, M., Carter, A.R., Shulman, G.L.: Common behavioral clusters and subcortical anatomy in stroke **85**(5), 927–941. doi:10.1016/j.neuron.2015.02.027
 39. Heitsch Laura, Ibanez Laura, Carrera Caty, Binkley Michael M., Strbian Daniel, Tatlisumak Turgut, Bustamante Alejandro, Ribó Marc, Molina Carlos, Dávalos Antoni, López-Cancio Elena, Muñoz-Narbona Lucia, Soriano-Tárraga Carol, Giralt-Steinhauer Eva, Obach Victor, Slowik Agnieszka, Pera Joanna, Lapicka-Bodzioch Katarzyna, Derbisz Justyna, Sobrino Tomás, Castillo José, Campos Francisco, Rodríguez-Castro Emilio, Arias-Rivas Susana, Segura Tomas, Serrano-Heras Gemma, Vives-Bauza Cristófol, Díaz-Navarro Rosa, Tur Silva, Jimenez Carmen, Martí-Fàbregas Joan, Delgado-Mederos Raquel, Arenillas Juan, Krupinski Jerzy, Culléll Natalia, Torres-Aguila Nuria P., Muiño Elena, Cárcel-Márquez Jara, Moniche Francisco, Cabezas Juan A., Ford Andria L., Dhar Rajat, Roquer Jaume, Khatri Pooja, Jiménez-Conde Jordi, Fernandez-Cadenas Israel, Montaner Joan, Rosand Jonathan, Cruchaga Carlos, Lee Jin-Moo, null null: Early neurological change after ischemic stroke is associated with 90-day outcome **52**(1), 132–141. doi:10.1161/STROKEAHA.119.028687. Publisher: American Heart Association. Accessed 2021-05-07
 40. NINDS, T.: The national institute of neurological disorders and stroke rt-pa stroke study group. tissue plasminogen activator for acute ischemic stroke. *N Engl J Med* **333**(24), 1581–7 (1995)