Individualizing deep dynamic models for psychological resilience data

Göran Köber^{1,2,*}, Shakoor Pooseh^{2,3}, Haakon Engen⁴, Andrea Chmitorz^{5,6,7}, Miriam Kampa^{5,8,9}, Anita Schick^{4,10}, Alexandra Sebastian⁶, Oliver Tüscher^{5,6}, Michèle Wessa^{5,11}, Kenneth S.L. Yuen^{4,5}, Henrik Walter^{12,13}, Raffael Kalisch^{4,5}, Jens Timmer^{2,3,14}, and Harald Binder^{1,2}

¹Institute of Medical Biometry and Statistics (IMBI), Medical Center—University of Freiburg, 79104, Germany

²Freiburg Center of Data Analysis and Modelling (FDM), University of Freiburg, Freiburg, 79104, Germany

³Institute of Physics, University of Freiburg, 79104, Germany

⁴Neuroimaging Center (NIC), Focus Program Translational Neuroscience (FTN), Johannes Gutenberg University Medical Center, Mainz, 55131, Germany

⁵Leibniz Institute for Resilience Research (LIR), Mainz, 55122, Germany

⁶Department of Psychiatry and Psychotherapy, Johannes Gutenberg University Medical Center, Mainz, 55131, Germany

⁷Faculty of Social Work, Health and Nursing, University of Applied Sciences Esslingen, Esslingen, 73728, Germany

⁸Department of Clinical Psychology, University of Siegen, 57076, Germany

⁹Bender Institute of Neuroimaging (BION), Department of Psychology, Justus Liebig University, Gießen, 35394, Germany

¹⁰Department of Public Mental Health, Central Institute of Mental Health, Medical Faculty Mannheim, Heidelberg University, Germany

¹¹Department of Clinical Psychology and Neuropsychology, Institute of Psychology, Johannes Gutenberg University, Mainz, 55131, Germany

¹²Research Division of Mind and Brain, Charité—Universitätsmedizin Berlin, Corporate Member of Freie Universität Berlin, Humboldt-Universität zu Berlin, and Berlin Institute of Health, Germany

¹³Berlin School of Mind and Brain, Humboldt-Universität zu Berlin, Germany

¹⁴CIBSS—Centre for Integrative Biological Signaling Studies, University of Freiburg, 79104, Germany

*koeber@imbi.uni-freiburg.de

ABSTRACT

Deep learning approaches can uncover complex patterns in data. In particular, variational autoencoders (VAEs) achieve this by a non-linear mapping of data into a low-dimensional latent space. Motivated by an application to psychological resilience in the Mainz Resilience Project (MARP), which features intermittent longitudinal measurements of stressors and mental health, we propose an approach for individualized, dynamic modeling in this latent space. Specifically, we utilize ordinary differential equations (ODEs) and develop a novel technique for obtaining person-specific ODE parameters even in settings with a rather small number of individuals and observations, incomplete data, and a differing number of observations per individual. This technique allows us to subsequently investigate individual reactions to stimuli, such as the mental health impact of stressors. A potentially large number of baseline characteristics can then be linked to this individual response by regularized regression, e.g., for identifying resilience factors. Thus, our new method provides a way of connecting different kinds of complex longitudinal and baseline measures via individualized, dynamic models. The promising results obtained in the exemplary resilience application indicate that our proposal for dynamic deep learning might also be more generally useful for other application domains.

Introduction

There currently is a renaissance of artificial intelligence techniques where, e.g., deep learning can provide human-level performance for detecting patterns in images, when trained on large data collections¹. Deep learning has further been successful with non-image data, where humans cannot directly detect patterns, and also for rather small sample sizes². Importantly, deep learning has recently been combined with time dynamic techniques such as ordinary differential equations (ODEs)^{3–6}, which promises both the advantages of hypothesis-driven modeling and detection of potentially complex patterns. This is why we turned to a combination of deep learning and ODEs when faced with modeling psychological resilience. We

propose a novel approach for combining variational autoencoders⁷ (VAEs)—which offer latent representations—with ODEs. Specifically, each individual receives an own set of ODE parameters for obtaining individualized models of their anticipated mental health response to stressor exposure and for subsequently identifying associated predictive factors from additional baseline measurements.

Psychological resilience is the maintenance or rapid recovery of a healthy mental state during and after times of adversity⁸. One element of the definition is that both mental health and adversity can change over time and may even do so permanently. Seminal resilience studies⁹ investigated how mental health changes in response to a single potentially-traumatic life event and identified common mental health trajectory patterns such as for resilient individuals (showing stably good or improving mental health in the months or years after the event) and more vulnerable individuals (showing stably poor or worsening mental health). These studies implicitly assume that the observed temporal changes in mental health are due to only *a single stressor event* and that individual differences in the mental health trajectory can be explained by some baseline individual characteristic. However, most individuals are *continuously exposed to more or less severe stressors*. These may include macrostressors (severe life events) but also more “mundane” microstressors, or daily hassles¹⁰, which are also known to have an impact on mental health^{11,12}. Further, major life events are likely to be followed by other macro- or microstressors^{13,14}. For instance, this means that deteriorations in mental health in the aftermath of a trauma may also have other causes than the trauma. Conversely, not developing long-term mental health problems may be a consequence of only moderate stressor exposure after the trauma. In this case, a resilient trajectory would not result from some protective baseline characteristics but benign life circumstances. The necessity to consider the ongoing, more or less continuous nature of stressor exposure is even more obvious when trying to understand resilience to longer-lasting or chronic adversity, such as adverse life circumstances or life transition phases. In short, investigating resilience should ideally involve repeated longitudinal measurements of stressors and mental health^{8,9,13,15}.

This necessity of longitudinal stressor and mental health measurements poses considerable challenges to data acquisition. Usually, mental health and stressor data are acquired intermittently, and not always are measurements equally spaced. The Mainz Resilience Project¹² (MARP) also faces these problems; self-report measurements are made using online questionnaires every three months $\pm$ two weeks. Additionally, a relevant number of observations (13.7%) are missing and the individual time series differ considerably in the numbers of observations (on average 9.1 observations, with a $\min = 2$, $\max = 14$, and $sd = 3.5$ in the data set used for learning the individual trajectories). Further, the three-monthly microstressor assessment of daily hassles only covers the past week. In this empirical context—which is prototypical for many longitudinal studies—the challenge is to transform the irregularly sampled mental health and stressor exposure data into continuous time series that are amenable to dynamic analysis.

In the current application, we learn the long-term individual individual mental health and stressor trajectories in the latent space, obtained from a VAE, with a system of ODEs that allows for encoding assumptions on the interplay of stressors and mental health. The ODE system is parameterized by another neural network, which receives longitudinal summary statistics as input, which also allows for different numbers of potentially irregularly spaced observation time points. On a more abstract level, our proposed new method showcases the versatility of combining deep learning techniques with modeling based on domain knowledge and underlines the usefulness of this particular new method in an application of practical relevance.

Dynamic techniques, such as differential equations, have already been proposed for modeling resilience¹⁶. Montpetit et al.¹⁷ presented a coupled version of the multivariate latent differential equation model¹⁸ for resilience, which estimates the actual values of the latent variable, as well as the slope and curvature simultaneously. However, this model does not allow for individualized parameters. Driver and Voelkle¹⁹ allow for individualized differential equations, but only for rather simple linear transformations for mapping from observed values to latent space, also requiring assumptions on the distribution of differential equation parameters, which may be problematic with respect to the topology in parameter space. In contrast, our proposal strives to combine the advantages of potentially complex non-linear transformations and individualized differential equations, making as little distributional assumptions as possible, as made possible by a combination of VAEs and individualized ODEs.

Naturally, there are other techniques that could potentially be used. For example, growth mixture models identify average starting levels and trajectories and allow for—normally distributed—deviation from the overall intercept and slope. However, these require an initial dimension reduction step (e.g., calculating a sum score or a factor solution), which is not informed by the longitudinal model and frequently based on rather simple linear transformations. By contrast, our method is designed to solve the problem of dimension reduction and trajectory modeling iteratively. Latent class mixture modeling has also been employed in resilience research⁹ to identify subgroups in the mental health response to trauma, to thereby derive categorical prediction targets (i.e., class membership); we expand on this by providing a set of continuous prediction targets. For enforcing smoothness across time, convolutional inference nets have been proposed, which also can find the lower-dimensional representations and connect them via Gaussian processes²⁰. Yet, this does not provide an explicit model, i.e., cannot be used for determining the anticipated stress response.

A brief sketch of the algorithm is provided next. After this, we show how these individualized models can be linked to baseline measurements for identifying resilience factors, before concluding with a discussion and outlook. In the methods

section, we provide more technical details of the proposed method, in particular on how the challenging estimation problem is addressed, and describe the MARP study on psychological resilience, which motivated methods development and is used for illustration. We also provide a Jupyter Notebook on [GitLab](#) that shows how our approach can recover structure from exemplary simulated data, and can be used as a starting point for adapting the approach to other applications.

Results

An overview of the algorithm, data structure, and training procedure

Figure 1 shows how the data flows through the algorithm. We map observations into a latent space using VAEs⁷. Specifically, separate VAE encoder networks, for mental health (mh; 28 General Health Questionnaire²¹ items, where higher values indicate more severe mental health problems) and stressor load (sl; the MIMIS battery of 58 daily hassles²²) respectively, reduce the number of dimensions to one for each type of measurement (Figure 1 a). The encoder weights, i.e., the parameters that determine the non-linear mapping, and the decoder weights, which specify the mapping from the latent space back to the level of original measurements, are trained by minimizing the evidence lower bound⁷ based on a Poisson log-likelihood to reflect the count character of the data.

The mapping of the observed values in the latent space is performed for each individual and time point where measurements are available, allowing to get a first impression of temporal trajectories by plotting the latent values along time (Figure 1 b). Subsequently, each individual time course is aggregated into several summary statistics, again separately for the latent mental health and stressor values. These summary statistics, e.g., include the integral over an interpolated step function and the difference of the first and last values (see Table 1 for a full list), and are chosen such that they do not require fixed observation time points or the same number of measurements for all individuals. These summary statistics then serve as inputs to a feed-forward ODEnet (Figure 1 c), which provides parameters for an ODE system. Thus, optimal combinations and transformations of the summary statistics are determined empirically. For the resilience application, we chose a rather simple ODE system, to reduce complexity when faced with a limited number of individuals and observations. Specifically, the parameters of the ODEnet are trained to minimize the squared distance between the solution of the ODE system and the latent values at the point in time when measurements were obtained from the individuals. We allow the stressor time course to additionally incorporate abrupt external inputs (i.e., sudden changes in the life situation). Accordingly, new observations can override the current state obtained from the ODE. For this purpose, the ODE solver is stopped at the measurement time points, and the latent values for stressors (but not for mental health) are updated. Subsequently, the individualized solutions of the ODE systems are visualized together with the latent values (Figure 1 d).

ODE system to model resilience trajectories in the latent space

Figure 1 d) shows an exemplary solution of an individualized ODE system, which connects expected values of the distributions in the latent space. We use an ODE system that couples the latent representations of mental health and stressor load. The exact design of such an ODE system is a crucial modeling decision since it governs how each component changes and, accordingly, requires domain expertise. Given our sample size and sampling frequency, we work with a rather simple ODE system to keep the number of estimated ODE parameters η_i small. Specifically, we use the system

$$\frac{df_{z_i^{mh}}}{dt} = \eta_{i,1}z_{i,t}^{mh} + \eta_{i,2}z_{i,t}^{sl} \quad (1)$$

$$\frac{df_{z_i^{sl}}}{dt} = \eta_{i,3}z_{i,t}^{sl} \quad (2)$$

where changes in $z_{i,t}^{mh}$, the latent value reflecting mental health problems, and $z_{i,t}^{sl}$, the latent value reflecting stressors load, are driven by the current own value. Additionally, mental health changes in response to stressor load. Separate parameter values $\eta_{i,1} - \eta_{i,3}$ for each individual i allow for individualized trajectories of mental health and stressors.

Negative values for $\eta_{i,1}$ and $\eta_{i,3}$ effectively realize system-inherent damping¹⁶, where high mental health and stressor values, respectively, are more quickly driven back to low latent mental health and stressor values. Thus, a high negative $\eta_{i,1}$ in particular reflects good recovery from mental health problems (since the majority of observations are mapped into the positive valued latent space or close to zero, see Figure 1 c). Positive values for $\eta_{i,2}$ realize the adverse effect of stressors on mental health. A low positive $\eta_{i,2}$ value thus reflects a low mental health impact of stressors in the individual.

At each realized measurement, the value of the integrator of the latent stressor load $f_{z_i^{sl}}$ is updated to the mapping of the actually observed value $z_{i,t}^{sl}$. This reflects the obvious notion that stressor levels are only partly driven by an endogenous property

Figure 1. The proposed new method has two essential parts, dimensionality reduction and individualized trajectory estimation. Both tasks are performed using neural networks (upper row). We train two VAEs—one for mental health (blue) and one for stressor load (red)—to estimate the distribution in the latent space for each observation. The variance of these distributions is expressed as size of the dots and reflects uncertainty. Summary statistics of the temporal pattern of latent values are used as inputs to a feed-forward neural network (the "ODEnet") which is trained to provide ODE parameters that minimize the squared distance of the ODE solution and the latent values, where latent stressor load values are updated at each measurement time point.

of the ODE system (i.e., damping) but mainly reflect exogenous forces, that is, the sudden occurrence or absence of stressors that lead to abrupt changes in the latent stressor values (see above).

The benefits of such an ODE system in comparison to discrete-time models like regression are crucial for analyzing the data from the MARP study at hand. Most importantly, differential equations take all available information at the precise time into account. Thereby, irregular sampling intervals and entirely (or partly) missing observations are dealt with straightforwardly.

Person-specific ODE parameters

An additional advantage of our approach is that the ODE parameters for every respondent are estimated in an unrestricted manner, i.e., without explicit constraints regarding their distribution, direction, or strength. This is accomplished by estimating the parameters of the ODE system $\eta_{i,1}$ to $\eta_{i,3}$ with the ODEnet—a feed-forward neural network—with 16 summary statistics as inputs (see Figure 1 c and Table 1 for details). While we are using linear ODE system for the specific resilience application, i.e., also an analytical solution exists, the proposed approaches enables more complicated non-linear systems. Importantly, the summary statistics are selected with data availability in mind. Since some MARP participants just started recently to provide data, all summary statistics need to be computable with only two observations of mental health and stressors (not necessarily at the same time), which is the minimum requirement to be included in the longitudinal analysis.

Input	Type	Scaled by
First obs of mh	O	1
First obs of sl	O	1
First obs mh – last obs mh	D	1
First obs mh – first obs sl	D	1
First obs mh – last obs sl	D	1
First obs sl – last obs sl	D	1
First obs sl – last obs mh	D	1
Last obs mh – last obs sl	D	1
Integral of mh	I	10
Integral of sl	I	10
Integral of mh ²	I	10
Integral of sl ²	I	10
Integral of mh (absolute value)	I	10
Integral of sl (absolute value)	I	10
Mean of autocorrelation mh	AC	100
Mean of autocorrelation sl	AC	100

Table 1. We include 16 summary statistics—gathered from the individual mappings of mental health (mh) and stressor load (sl) to the latent space (see Figure 1 c)—as inputs to the ODEnet. We additionally scale them according to their range and type (Observation, Difference, Integral, and AutoCorrelation) for achieving numerically similar inputs.

The ODEnet provides the individual parameters of the ODE system as outputs. Accordingly, the overall structure of how mental health and stressors change is the same for all respondents and specified by the ODE system. Yet, the values of the summary statistics—included in the ODEnet—differ from person to person. Therefore, each individual receives an own variant of the ODE system. The ODEnet is trained to minimize the squared difference between the ODE solutions $f_{z_i^{mh}}$ and $f_{z_i^{sl}}$ and the latent values $z_{i,t}^{mh}$ and $z_{i,t}^{sl}$, using a standard neural network backpropagation algorithm. In contrast to previous methods (e.g., growth mixture models), the individual parameters do not obey any predefined distribution; they can be estimated in the absence of such restrictions by optimizing the fit to the data.

Quantification and prediction of resilience

In addition to longitudinal measurements of mental health and stressor load, the MARP study incorporates a large battery of additional baseline measurements for identifying potential resilience factors, i.e., characteristics of individuals that might predict reduced reactivity of mental health to stressor load. For identifying such resilience factors, we can either predict the individual parameters η_i directly or apply an artificial stress test to each individual (see Figure 2). This artificial stress test means that the latent mental health value of an individual is initialized to an average value, and the latent stressor load value is set to an initial value that corresponds to a high stressor level. The stress reaction pattern of each individual then is characterized by how the latent values develop according to the individualized ODE system, if no further external input is provided. Predicting each of the ODE parameters separately allows for investigating the different dimensions of resilience in our model (i.e., recovery

Figure 2. The artificial stress test induces a considerable amount of stress ($z_{i,0}^s = 2$) on each individualized ODE system and captures how the mental health (blue) and stressor load (red) are predicted to develop in the latent space (y-axis) over time (x-axis). Prediction targets are derived from the values of mental health at predefined time points (ptp 1-4) of 5, 9, 15 and 20 months.

Figure 3. Inclusion frequencies of potential resilience factors in % (shades of gray) using a cross-validated lasso analysis ($n = 155$; $p = 350$; 1000 repetitions) predicting mental health at predefined time points (ptp) of the artificial stress test, i.e., the intersection of the blue lines with the vertical, dotted lines in Figure 2). We identified potential resilience factors in all three data modalities (i.e., questionnaires, behavioral measures, and proteomics).

of mental health ($\eta_{i,1}$) and reactivity of mental health to stressors ($\eta_{i,2}$) independently. In contrast, the artificial stress test considers all parameters jointly, to obtain a holistic picture of resilience. This also serves as a diagnostic tool to investigate whether the training of the algorithm succeeded. We choose four points in time (after 5, 9, 15, and 20 months) and predict the individual mental health values with measurements from the baseline battery, covering proteomics, behavioral measures, and questionnaires.

Specifically, we use the lasso, a regularized regression approach that provides variable selection²³. The amount of regularization is chosen according to prediction performance, specifically by 6-fold cross-validation. For a robust result, this procedure is repeated 1000 times on resampling data sets, obtained by sampling with replacement, resulting in inclusion frequencies²⁴. Thus, a baseline measurement would receive inclusion frequency equal to or close to zero, if there would be no connection between the baseline measurement and our proposed quantification of resilience by the stress test. Quite in contrast, we find rather strong relations between our derived quantification of resilience and some baseline measurement, which could thus be considered to be potential resilience factors (see Figure 3). Since the present paper has its focus on the proposed new method, and not on resilience research, we do not report specific selected variables and refrain from further interpretation.

Discussion

We combine deep generative learning with a system of ODEs to obtain individualized dynamic models, and illustrate this novel approach in an application on psychological resilience. Our proposal methods can handle problems typical for real-world applications, specifically irregular measurement time points and limited sample size. In particular, the augmentation of neural networks with ODEs allows us to incorporate domain expertise, which can dramatically reduce the sample size requirements of deep learning⁶. An important ingredient of this approach is differential programming, i.e., automatic differentiation in complex models, which allows for tailoring algorithms to particular needs, such as combining neural networks with ODEs. This framework is more specifically used in our proposed method to allow the integrator of the ODE solver to jump to updates of the latent stressor load values.

We show that this flexible, non-linear, and versatile combination of scientific modeling and machine learning can be utilized to quantify resilience and identify resilience factors, and agree with others^{3,6} that neural differential equations promise considerable opportunities for this particular field and other disciplines working with longitudinal data.

From a subject matter point of view, the most substantial benefit of our approach are the resulting individualized dynamic models. Building upon the overall structure of stress research, according to which mental health problems change due to stressor load, and incorporate the main lesson from resilience research that there are significant individual differences in how people deal with stressors and recover after stressful episodes. In fact, these individual differences are what resilience research is investigating. Instead of reporting population-level prevalence (like, e.g., what percentage of people can be labeled as resilient) we are interested in modeling the mental health trajectory of each person in response to his or her stressor load. This perspective on each individual also addresses the needs of health professionals who are interested in personalized predictions tailored to their particular client, strongly informed by his or her unique history. Hence, we are interested in the most flexible approach that allows for capturing personal nuances best. Compared to this, we argue that it is of secondary interest (at most) to assign people to categorical labels like resilient and vulnerable or force individual deviations from the grand mean to be aligned with predefined distributions. In this regard, we argue that the continuous individual differences—expressed here as ODE parameters and harnessed with an artificial stress test—are meaningful quantifications of resilience and help to identify resilience factors from large baseline batteries.

A limitation of the proposed model is the simplicity of the ODE system. There are valid theoretical reasons to assume that resilience might be better captured using an additional second-order derivative (i.e., acceleration) of mental health problems. While our first-order ODE system was a deliberative decision in consideration of the small number of observations and individuals (for dynamic modeling and deep learning standards at least), future research might experiment with more complicated systems of differential equations. Also, our approach assumes stationarity and accordingly neglects the possibility of long-term learning and plasticity (i.e., temporal changes in resilience). While we could extend our model to weaken this assumption, our main focus in this work is on the identification of resilience factors based on all available longitudinal measures. Future work will investigate repeated baseline observations and time-dependent predictors. Yet, already the individualized dynamic models enabled by our new method provide an important building block for better understanding psychological resilience. Therefore, we are confident that other applications that required individualized dynamic models might benefit from our approach.

Methods

Dimensionality reduction per time point with a VAE for count data

We choose variational autoencoders (VAEs) to find a lower-dimensional latent representation because they provide a flexible framework with numerous potential extensions and applications. VAEs comprise a recognition model and a generative model. The purpose of training the recognition model (aka inference model or encoder) is to find the variational parameters ϕ of a neural network to approximate the posterior distribution of the latent variable z given the inputs x , i.e., $q_\phi(z|x)$. The parameters of the generative model (aka decoder) θ are trained to increase the log-likelihood of the inputs given random samples from the posterior distributions. To train the recognition and generative model simultaneously, we maximize the evidence lower bound (ELBO) of the marginal log-likelihood

$$\begin{aligned} \log p(x) &= \sum_{n=1}^N \sum_{t=1}^T \log p(x) \\ &\geq \sum_{n=1}^N \sum_{t=1}^T \mathbb{E}_{q(z|x)} \left[\sum_{p=1}^P \log p(x|z) \right] - \sum_{n=1}^N \sum_{t=1}^T \text{KL} (q(z|x) || p(z)) \end{aligned} \quad (3)$$

where the first term of the right hand is the expectation of the log-likelihood of x given z with respect to $q(z|x)$. The

Kullback-Leibler divergence (D_{KL}) penalizes deviations of the posterior from the prior. For computation, we plug in the Poisson log-likelihood and the closed form of D_{KL} for a Gaussian prior and posterior. Thereby, our training objective becomes

$$\mathcal{L}(\theta, \phi; x) = \underbrace{\frac{1}{2} \sum_{i=1}^N (1 + \log((\sigma_i(x))^2) - (\mu_i(x))^2 - (\sigma_i(x))^2)}_{D_{KL}(\phi)} - \underbrace{\sum_{i=1}^N \lambda - x \times \log(\lambda)}_{\text{Poisson reconstruction error } (\theta)} \quad (4)$$

where μ_i and σ_i are the mean and standard deviation of the observation and λ is the expected value of the Poisson distribution. The Poisson distribution does not perfectly fit the data since we found moderate levels of overdispersion for some of our items. Yet, the Poisson distribution avoids an additional parameter, which might be difficult to determine, by coupling expected value and dispersion. We used J nodes and $\tanh$ activation functions in the middle layers. In the final layer, we used a $ReLU$ activation function to strictly pass non-negative values to λ . All neural networks were trained with Flux.jl²⁵ and the Adam²⁶ optimizer with standard decay rates ($\beta_1 = 0.9$ and $\beta_2 = 0.999$).

Individual ODE parameter estimation with the ODEnet

The parameters of the ODEnet τ were trained to minimize squared difference of the trajectory at the precise point in time t and the mean of the latent space distribution $\mu_{z_{i,t}}$.

$$\mathcal{L}_{ODE}(\tau, \eta_i) = \sum_{n=1}^N \sum_{t=1}^T (f_{z_i}(t, \eta_i) - \mu_{z_{i,t}})^2 \quad (5)$$

The ODEnet is a separate feed-forward neural network with two layers and 16 inputs; the mid-layer and end-layer have three nodes (due to data size considerations). As inputs, we use a mixture of integrals, first and last observations, differences, and autocorrelations (see Table 1).

We used a piecewise constant step function to approximate the integral of the individual trajectories defined as $s_i = \sum_{m=1}^{T_i} c_m l(I_m)$ where c_m is the value of $z_{i,t}^{mh}$ or $z_{i,t}^{sl}$ on the interval $I_m = (t_{i-1}, t_i)$ and $l(I_m)$ the length ($t_i - t_{i-1}$) of this interval. We also included s_i^2 . It was left to gradient descent/backpropagation to find the best combination of these inputs to minimize $\mathcal{L}_{ODE}(\tau, \eta_i)$ given the value of the ODE at a particular point in time $f_{z_i}(t_i)$ which, in turn, depends on the individual set of ODE parameters η_i .

We use a $ReLU$ activation function in the middle layer and no transformation in the final layer. The choice of inputs is restricted by the minimum requirement of two observed values for each group of variables at any point in time. We scale the inputs to ensure approximately equal numerical size. For reasons explained above, we update the integrator at every realized measure of $z_{i,t}^{sl}$. Such steps are implemented in DifferentialEquations.jl²⁷ and differentiable through neural nets via DiffEqFlux.jl²⁸. All solved ODE systems start the the first measurement $\mu_{z_{i,1}}$. The ODEnet was trained with the Adam optimizer²⁶ and standard decay rates ($\beta_1 = 0.9$ and $\beta_2 = 0.999$). To deal with unit non-response, $\mathcal{L}(\theta, \phi; x)$ and $\mathcal{L}_{ODE}(\tau, \eta_i)$ are only evaluated at actual measurement time points.

Training

We experimented with different training strategies. The best results are obtained with a combination of separate pretraining of the VAE and ODEnet. Subsequently, we also compared joint training of all involved components vs. a sequential strategy where we trained the VAE components separate from the ODEnet in one epoch. This leads to further decreases in the loss components. In this analysis, we pretrain both VAEs with a learning rate $\alpha = 5^{-5}$ for 80 epochs. Then, we pretrain the ODEnet separately for 50 epochs with $\alpha = 1^{-6}$. Then, we sequentially train both nets in the same epoch with $\alpha = 1^{-6}$ for 100 epochs.

MARP Data

The Mainz Resilience Project (MARF) is an ongoing study that started in 2016 with a planned study duration per participant of seven years and is conducted by the University Medical Center Mainz and the Leibniz Institute for Resilience Research¹². All available observations ($n_{sl} = 1.233$ and $n_{mh} = 1.196$) were used to train the VAEs. To be included in the longitudinal analysis, at least two observations of mental health and stressor load—at any point in time within the observation window of 3.5 years—were necessary. This basic requirement was met by $N = 166$ participants. Four respondents were excluded due to unusual patterns. In this longitudinal data set, on average 9.1 observations ($min = 2$, $max = 14$, $sd = 3.5$) per person are available. The majority of the longitudinal measures are gathered via an online assessment roughly every three months. Besides the request to provide data on their mental health and stressors (daily hassles and life events, of which only daily hassles are used here), the participants went through an extensive baseline examination, which includes in total $p = 350$ neuroimaging,

behavioral, proteomics, and survey data variables. One hundred fifty-five participants met the longitudinal requirements and had a complete baseline battery; hence, they were included in the lasso analysis (see Figure 2). Participants were recruited in a critical life phase aged 18 to 20 at study inclusion with a prehistory of critical life events.

References

1. Esteva, A. *et al.* Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115 (2017).
2. Hess, M., Lenz, S., Blätte, T. J., Bullinger, L. & Binder, H. Partitioned learning of deep Boltzmann machines for SNP data. *Bioinformatics* **33**, 3173–3180, DOI: [10.1093/bioinformatics/btx408](https://doi.org/10.1093/bioinformatics/btx408) (2017).
3. Chen, T. Q., Rubanova, Y., Bettencourt, J. & Duvenaud, D. K. Neural ordinary differential equations. In *NeurIPS* (2018).
4. De Brouwer, E., Simm, J., Arany, A. & Moreau, Y. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. In *Advances in Neural Information Processing Systems*, 7377–7388 (2019).
5. Yıldız, Ç., Heinonen, M. & Lähdesmäki, H. ODE²VAE: Deep generative second order ODEs with bayesian neural networks. *arXiv preprint arXiv:1905.10994* (2019).
6. Rackauckas, C. *et al.* Universal differential equations for scientific machine learning (2020).
7. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv* (2013).
8. Kalisch, R. *et al.* The resilience framework as a strategy to combat stress-related disorders. *Nat. Hum. Behav.* **1**, 784 (2017).
9. Bonanno, G. A., Westphal, M. & Mancini, A. D. Resilience to loss and potential trauma. *Annu. Rev. Clin. Psychol.* **7**, 511–535, DOI: [10.1146/annurev-clinpsy-032210-104526](https://doi.org/10.1146/annurev-clinpsy-032210-104526) (2011).
10. Hahn, S. E. & Smith, C. S. Daily hassles and chronic stressors: Conceptual and measurement issues. *Stress. Medicine* **15**, 89–101, DOI: [10.1002/\(SICI\)1099-1700\(199904\)15:2<89::AID-SMI789>3.0.CO;2-O](https://doi.org/10.1002/(SICI)1099-1700(199904)15:2<89::AID-SMI789>3.0.CO;2-O) (1999).
11. Serido, J., Almeida, D. M. & Wethington, E. Chronic stressors and daily hassles: Unique and interactive relationships with psychological distress. *J. Heal. Soc. Behav.* **45**, 17–33 (2004).
12. Kalisch, R. *et al.* A generic solution for the operationalization and measurement of resilience and resilience processes in longitudinal observations: rationale and basic design of the MARP and LORA studies. *PsyArXiv* DOI: [10.31234/osf.io/jg238](https://doi.org/10.31234/osf.io/jg238) (2020).
13. Kalisch, R., Müller, M. B. & Tüscher, O. A conceptual framework for the neurobiological study of resilience. *Behav. Brain Sci.* **38**, DOI: [10.1017/S0140525X1400082X](https://doi.org/10.1017/S0140525X1400082X) (2015).
14. Norris, F. H. & Elrod, C. L. Psychosocial consequences of disaster. *Methods for disaster mental health research* 20–42 (2006).
15. Kalisch, R. *et al.* Deconstructing and reconstructing resilience: a dynamic network approach. *Perspectives on Psychol. Sci.* **14**, 765–777 (2019).
16. Boker, S. M., Montpetit, M. A., Hunter, M. D. & Bergeman, C. S. Modeling resilience with differential equations. In *Individual pathways of change: Statistical models for analyzing learning and development*, 183–206 (APA, 2010).
17. Montpetit, M. A., Bergeman, C. S., Deboeck, P. R., Tiberio, S. S. & Boker, S. M. Resilience-as-process: negative affect, stress, and coupled dynamical systems. *Psychol. Aging* **25**, 631–640 (2010).
18. Boker, S., Neale, M. & Rausch, J. Latent differential equation modeling with multivariate multi-occasion indicators. In *Recent developments on structural equation models*, 151–174 (Springer, 2004).
19. Driver, C. C. & Voelkle, M. C. Hierarchical Bayesian continuous time dynamic modeling. *Psychol. Methods* **23**, 774 (2018).
20. Fortuin, V., Baranchuk, D., Rätsch, G. & Mandt, S. GP-VAE: Deep probabilistic time series imputation (2019).
21. Goldberg, D. P. *et al.* The validity of two versions of the GHQ in the WHO study of mental illness in general health care. *Psychol. Medicine* **27**, 191–197 (1997).
22. Chmitorz, A. *et al.* Assessment of microstressors in adults: Questionnaire development and ecological validation of the main inventory of microstressors. *JMIR Ment Heal.* **7**, e14566, DOI: [10.2196/14566](https://doi.org/10.2196/14566) (2020).
23. Hastie, T., Tibshirani, R. & Wainwright, M. *Statistical learning with sparsity: the lasso and generalizations* (Chapman and Hall/CRC, 2015).

24. Sauerbrei, W., Buchholz, A., Boulesteix, A.-L. & Binder, H. On stability issues in deriving multivariable regression models. *Biom. J.* **57**, 531–555 (2015).
25. Innes, M. *et al.* Fashionable modelling with flux. *CoRR* (2018).
26. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv* (2014).
27. Rackauckas, C. & Nie, Q. DifferentialEquations.jl – A performant and feature-rich ecosystem for solving differential equations in julia. *J. Open Res. Softw.* **5**, DOI: [10.5334/jors.151](https://doi.org/10.5334/jors.151) (2017).
28. Rackauckas, C. *et al.* Diffeqflux.jl – A julia library for neural differential equations. *CoRR* (2019). [1902.02376](https://arxiv.org/abs/1902.02376).

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 777084 (G.K., S.P., H.E., O.T., R.K.) and the State of Rhineland-Palatinate, Germany (MARF program, DRZ program, Leibniz Institute for Resilience Research; M.K., A.Sch., M.W., O.T., R.K.).

Author contributions statement

H.B. and G.K. developed the idea, constructed and tested the models, and wrote the paper. All authors contributed through discussions and reviewed and approved the manuscript.

Additional information

Competing interests: R.K. receives advisory honoraria from JoyVentures, Herzlia, Israel. We have no conflict of interest to declare.