

Using Variational Autoencoder to Develop and Validate a Compact, Deep Representation of Digital Clock Drawing Test for Classifying Dementia

Sabyasachi Bandyopadhyay

University of Florida

Catherine Dion

University of Florida

David J. Libon

Rowan University

Patrick Tighe

University of Florida

Catherine Price

University of Florida

Parisa Rashidi (✉ parisa.rashidi@bme.ufl.edu)

University of Florida

Research Article

Keywords: Clock Drawing Test (CDT), Variational Autoencoder, Deep Representation, Classifying Dementia

Posted Date: January 11th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1207133/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Scientific Reports on May 14th, 2022. See the published version at <https://doi.org/10.1038/s41598-022-12024-8>.

Abstract

The Clock Drawing Test (CDT) is an inexpensive tool to screen for dementia. In this study, we examined if a semi-supervised deep learning (DL) system using Variational Autoencoder (VAE) can extract atypical clock features from a large dataset of unannotated CDTs (n=13,580) and use them to classify dementia (n=18) from non-dementia (n=20) peers. The classification model built with VAE latent space features adequately classified dementia from non-dementia (0.78 Area Under Receiver Operating Characteristics (AUROC)). The VAE-identified atypical clock features were then reviewed by domain experts and compared with existing literature on clock drawing errors. This study shows that a semi-supervised deep learning (DL) analysis of the CDT can extract important clock drawing anomalies that are predictive of dementia.

Introduction

The Clock Drawing Test provides a method for screening individuals for dementia. The clock drawing test consists of two parts: the command condition, where participants are required to “draw the face of a clock, put in all the numbers, and set the hands for *ten after eleven*”; followed by the copy test condition where participants are instructed to copy a model clock. Accurate clock drawing relies on various cognitive domains, and subtle changes in drawing can provide rich information about underlying cognitive functioning^{1,2}. Command-based clock drawing requires comprehension for instructions, recalling the semantic attributes of a clock, working memory to processing the linguistic components of the test instructions, effective mental planning, visuospatial processing, and motor skills to execute the drawing.³ Although the copy condition also requires an array of cognitive function capabilities, it mainly relies on efficient visual scanning abilities, visuoconstruction, and executive functioning.⁴⁻⁶ Literature also suggests that performance on the clock drawing test correlates with the Mini-Mental State Examination (MMSE) total score which is an alternative test of cognitive impairment.^{7,8}

Several analog clock scoring systems have been proposed in the past.^{1,9,10} These scoring systems range from nominal (right/wrong) to elaborate 22- or 31-point scoring.¹¹ Some are based on analysis of errors assessing semantics, graphomotor functioning, and executive control.¹ Different scoring protocols have been shown to have similar psychometric properties.¹² However, Spenciere et al. (2017) suggest that the subjective human component for interpreting clock drawing performance can yield different outcomes¹³. Furthermore, Price and colleagues found variance in intra and inter-rater reliability^{14,15}. Variability in rater scoring introduces ambiguities that can potentially negatively impact the robustness of any diagnostic test based on the CDT.

Deep Learning (DL) models could obviate this problem due to their ability to automatically extract high-level features from the data without any *a priori* feature engineering. Such high-level features are extracted in a data-driven manner by continuously assessing correlations between simpler features. The generality and predictive power of this nested hierarchy of features is only limited by the size of training

data. DL models (given they have sufficient data to train), therefore, present an opportunity for developing objective scoring criteria for more robust clinical decision-making.

In this study, we developed an interpretable DL model to automatically learn key clock drawing features, and then validated this learning with “disease” classification in a sample of clinically diagnosed individuals with Alzheimer’s Disease (AD) or Vascular dementia (VaD) versus non-dementia peers. We used the final image output of the digital clock drawing test (dCDT)¹⁶ for this work. The dCDT uses digital pen technology with associated smart paper, which can record every pen stroke during drawing, allowing researchers to examine a multitude of clock drawing elements like latencies between pen strokes, graphomotor elements such as the size of the clock face and, the total number of pen strokes, etc.^{5,17} This technology, therefore provides additional benefits compared to the traditional clock drawing test because it can analyze the process by which the drawing was created rather than solely relying on the final created image. Automated scoring models developed on the dCDT in recent years^{18–20} have significantly outperformed traditional pen and paper clock drawing test. Their AUROC scores ranging from 0.89 to 0.93, compared against published AUROC scores from existing clinician scoring systems of 0.66–0.79.²¹ In this study the protocol of kinematic, time-based or latency, and visuospatial features extracted by the dCDT were not used. Rather, we used only the clock image or actual drawing produced by patients and research participants to train a DL model to screen dementia.

Despite its promise, many disadvantages plague the successful application of DL pipelines to medical image analysis. First, the number of parameters in powerful DL models are many orders of magnitude higher than in standard machine learning models such as logistic regression. Generally, these parameters are the network weights of a Deep Neural Network (DNN), which allow the network to model complex, arbitrary input-output relationships. The use of a large parameter space begets the need for large, labeled datasets to train these DNNs. A DNN uses the error values defined as the distance between the predicted class probabilities and the ground truth to iteratively update its parameters. As a result, the final parameter values become entirely dependent on the dataset and typically fail to generalize to other classification tasks. In this work, due to availability of a large unlabeled dataset and a significantly smaller labeled dataset, we decided to use a semi-supervised DL model that can learn the intrinsic variations in clock images from the unlabeled data and use the small labeled data to solve the classification task with minimal fine-tuning.

We used a Variational Autoencoder (VAE) model to solve the self-supervised learning task. VAE is a generative model which aims to learn a joint probability distribution over all variables present in a dataset.^{22,23} This technique uses the accurate reconstruction of input images as an objective to learn a low-dimensional latent representation in the form of a pre-defined prior distribution. Deep generative models have been shown to improve classification accuracy in semi-supervised learning settings, especially when one has few labeled examples and many more unlabeled examples.^{24,25} We used a large unlabeled dataset of clock drawings to train the VAE and a considerably smaller labeled dataset to subsequently fine-tune the trained VAE encoder network. The encoder represents the part of the VAE

network that encodes a clock drawing into a low-dimensional latent space. This research is the first attempt at using semi-supervised DL models to analyze clock drawings. The primary aim of this project is to demonstrate the ability of semi-supervised deep learning to build an efficient classifier capable of distinguishing individuals with AD or VaD from control participants despite the lack of large quantities of labeled data necessary for training traditional machine learning methods.

Results

Participants

Table 1 describes the pertinent demographic characteristics for the participants who completed clock drawings for the training and classification datasets. Individuals in both datasets were above the age of 60. All subjects in the classification cohort completed the clock drawing test to both conditions. In the training dataset 3 subjects could not draw the command clocks. Within the classification cohort, the dementia subgroup was significantly older, less educated and had lower MMSE total score on average compared to their non-dementia peers. In comparison to the training cohort, the classification cohort was predominantly comprised of white female participants.

Table 1
Demographics of Cohorts.

Dataset	Number of samples	Mean Age (S.D)	Mean Education (S.D)	% of Female	% of Caucasian	Mean MMSE total score	Mean MoCA total score
Training clocks	13,580	73 (6)	13 (3)	51	85	26 (4.0)	N/A
Dementia	71	80 (6)	13 (3)	68	100	22 (2.6)	N/A
Controls	80	68 (6)	16 (2)	67	94	29 (1.1) *	25 (2.2) †
Abbreviations. S.D, Standard Deviation. MMSE, Mini-Mental State Examination. MoCA, Montreal Cognitive Assessment. *This value is available only for 24 subjects in the Controls cohort. † This value is available only for 56 subjects in the Control cohort.							

Classification dataset (fine-tuning cohort)

Figure 1A illustrates the reconstructions of clock drawings performed by the VAE as a function of its latent variables. The smooth transition between the reconstructed clocks observed here results from the normal distribution to which the VAE latent variables are restricted. Figure 1B shows the distribution of

reconstructions of dementia/control clocks from the fine-tuning dataset as a function of their latent space vectors.

The scatterplot shows visible separation between dementia and control clocks. The reconstructions denote features which are salient for human perception such as digits and ticks are not captured by the VAE latent space. Instead, the VAE latent space captures statistically features of clock drawings such as eccentricity, size, size of clock hands and distance of clock hands from the geometric center (Fig. 2). These features are correlated to the two latent dimensions of the VAE latent space and are mutually entangled in this space.

The latent manifold upon which the trained VAE projects clock drawings is a two-dimensional vector space that can be functionally divided into five regions which delineate the variation of different clock drawing features and anomalies. Figure 2 shows a traversal over these five regions. The VAE had no prior information about the generative features of clock drawing. Therefore, the clock features appear in an entangled/mutually correlated manner in the two-dimensional latent space of the VAE.

Latent Dimensions

The first latent dimension (plotted along the X-axis) is abbreviated as Z_0 , and the second latent dimension (plotted along the Y-axis) is abbreviated as Z_1 . The left half of the latent manifold space ($Z_0 < 0$) is concerned with the direction of the eccentricity of a clock drawing (Fig. 2, Supplementary Video V1). The eccentricity reverses from left to right as Z_1 given Z_0 ($Z_1|Z_0$) changes from -4 to +4, passing through a point of zero eccentricity (circular clock at $Z_1|Z_0 = 0$). The size of the clock drawing is a correlated variable which decreases as $Z_1|Z_0$ increases from -4 to +4 in the left half of the two-dimensional manifold. On the other hand, the right half of the latent manifold ($Z_0 > 0$) is related to the distance of the point of intersection of clock hands from the clock's geometric center (Fig. 2, Supplementary Video V2).

The point of intersection of clock hands moves downwards from the geometric center as $Z_1|Z_0$ traverses from -4 to +4 in this region. This change is also associated with a loss of the circular periphery of the clock which is an important anomaly present in clocks drawn by patients with advanced stages of dementia. The top of the latent space ($Z_1 > 0$) encodes an increasing length of clock hands; and distance of point of intersection of clock hands to geometric center mixed with each other (Fig. 2, Supplementary Video V3). Length of clock hands and area of the clock face increase as $Z_0|Z_1$ changes from -4 to +4. The bottom half of the latent space ($Z_1 < 0$) encodes the eccentricity of the clockface (Fig. 2, Supplementary Video V4). Eccentricity decreases as $Z_0|Z_1$ changes from -4 to +4 in this region of the latent space. Furthermore, this feature is interlinked with an increase in clockface area as $Z_0|Z_1$ changes from -4 to +4 in this region. Finally, the X-axis which traces the change in $Z_0|(Z_1=0)$ purely encodes the size of the clockface, evident from the increasing clarity of the clock drawing along this line (Fig. 2, Supplementary Video V5).

This analysis reveals that many physically understandable clock features and anomalies are encoded in different regions of the latent space of the trained VAE. It also shows that in many cases multiple

physical features are entangled with each other.

Classification Dataset (testing cohort)

The trained VAE encoder was used to perform classification after fine-tuning for 10 epochs over the fine-tuning dataset described previously. Table 2 shows the performance of the semi-supervised classifier constructed by employing the trained encoder network from the VAE. 95% confidence intervals denote the performance of the model over 100 bootstrapped examples of the test data. Figure 3 demonstrates the Receiver Operating Characteristic (ROC) curve and the Precision-Recall (PR) curve obtained upon classification on test data using this network.

Table 2
Performance of semi-supervised classifier on test data

Task	AUROC (95% C.I.)	Accuracy (95% C.I.)	F1- score (95% C.I.)	Precision (95% C.I.)	Sensitivity (95% C.I.)	Specificity (95% C.I.)	NPV (95% C.I.)
Dementia prediction	0.78 (0.59 – 0.91)	0.71 (0.54 – 0.83)	0.70 (0.51 – 0.82)	0.71 (0.47 – 0.87)	0.69 (0.49 – 0.86)	0.74 (0.54 – 0.87)	0.74 (0.49 – 0.88)
Abbreviations. AUROC, Area Under the Receiver Operating Curve; C.I, Confidence Interval; NPV, Negative Predictive Value.							

Discussion

VAE trained on reconstructing digital clock drawings compress the relevant information present in a clock image into a highly informative 2-dimensional vector. The result appears to have useful properties for classifying dementia versus non-dementia. The latent space of the VAE can be used to generate artificial clock drawings which show statistical resemblance to the human-drawn clock dataset, but do not replicate salient features/details of clock such as digits, hands, and ticks which are central to scoring clocks using traditional scoring techniques. Instead of local features, the trained VAE latent space captures global features such as clockface eccentricity, clockface area, length of hands and distance from the point of connection of hands to the clock center. Some of these features have been separately identified by domain experts as pertinent in distinguishing amongst various subtypes of dementia, as well as in separate cognitive tasks. For example, smaller clock face area is associated with micrographia and subcortical disease profiles where there is presence of primary executive dysfunction (e.g., Parkinson’s disease)⁹. Individuals with executive dysfunction and Parkinson’s disease also exhibit planning deficits in laying out numbers⁹. Clock face hand placement is also a consideration for disinhibition and visual attention difficulties¹. The obfuscation of fine details of the clock drawing might be an inevitable trade-off related to the preprocessing step where clock drawings are resized to 100X100 to input them to the VAE. Future studies in this domain can find similar relationships between obliqueness of clock image or distance of the point of contact of clock hands from the center with underlying pathologies.

Furthermore, this analysis reveals that many physically understandable clock features and anomalies are encoded in different regions of the VAE latent space in a mutually entangled manner. In future work, we will investigate if disentangling these generative features improves the performance in dementia classification. To do this, we will use latent variable construction using advanced VAE models such as Factor-VAE²⁶⁻²⁸. In addition, we will investigate if a classifier built using the generative factors discovered in this paper can distinguish dementia from control. These generative factors are recorded in the software associated with the digital version of the clock drawing test that is used in this research¹⁷. We will augment these features with other graphomotor, and latency features solely available through the dCDT to examine the relative importance of these features in classifying dementia from nondementia.

The dCDT uses digital pen technology with associated smart paper, which record each pen stroke allowing researchers to access an array of variables inherent in the process by which the clock drawing was made in addition to the drawing itself. The dCDT can also capture full videos of the clock drawing test which can be analyzed using state of the art sequential deep learning models such as Vision Transformers²⁹ to build high-resolution screening tools for dementia. These future research works will allow us to explore the wide range of data collection capabilities of the dCDT.

Importantly this methodology presents an important advance in bidirectional translational neuroscience involving AI. Here, we have used mechanistic understanding of dCDT tasks, developed in concert with neuroimaging studies, to train the latent representations of the VAE as part of a series of forward-translational experiments. Based upon visual inspection of the latent representations of the VAE, and in concert with the classification results, domain experts can use these findings to identify novel feature combinations of the CDT images and map these to gold standard cognitive assessments and/or neuroimaging findings. This bidirectional translational opportunity emphasizes the importance of methods which are sensitive to domain-level concerns including interpretability and mechanistic grounding.

This study is not without its limitations. Firstly, the predictive model needs to be validated on a larger dataset as the current validation dataset consists of 18 dementia and 20 non-dementia subjects. Secondly, the dementia group is significantly older and less educated than the non-dementia group, thereby making the comparison unmatched. Furthermore, the objective of classification is more general than in other studies in this domain. In the future, we plan to use the current methodology to differentiate between different types of dementia such as AD, VaD, mild cognitive impairment (MCI), amnesic-MCI, dysexecutive-MCI and Parkinson's disease dementia. The encoding capability of the VAE was shown to depend on the amount of white space present in the sketches (Supplementary Fig S1). This also limited our study as we had to discard the clock drawing images which were larger than 40,000 pixels in size.

In summary, this study demonstrated that the traditional pen and paper clock drawing test can inform contemporary state-of-the-art DL models to classify participants diagnosed clinically with dementia. Since our classifier solely relies on the clock drawing and is trained in an unsupervised manner, it can potentially leverage large publicly available clock drawing datasets for creating similar classifiers with

minimal fine-tuning for other conditions such as delirium and traumatic head injury. It can also be used to compare clock drawing images before and after a surgery or other insults. The same VAE encoder network, with minimal fine-tuning, can create new classifiers for different tasks in a short amount of time while using scant computational resources. Furthermore, such models can also be trained on the fly to monitor clock drawing changes over time, thereby serving as a simple, effective monitoring system of the neuro-cognitive health of a subject.

Conclusion

In summary, in this study, we have shown that semi-supervised DL models trained on unlabeled clock datasets have the potential of extracting generative clock features which are sufficiently informative to construct a dementia classifier. This is the first study using semi-supervised DL methods to analyze digital clock drawings. In the future, we will expand the classifier using graphomotor and latency features available from the dCDT and improve the latent space of the VAE using disentangled VAE in search for mapping between specific clock drawing errors and underlying cognitive pathologies.

Methods

Participants

Study materials were collected from digital clock drawing consortium data between the University of Florida (UF) and New Jersey Institute for Successful Aging (NJISA), Memory Assessment Program, School of Osteopathic Medicine, Rowan University. The Institutional Review Boards (University of Florida IRB and Rowan University, New Jersey Institute for Successful Aging IRB) approved this investigation, and permission was collected from the participants at both institutions through their signatures on informed consent forms. All study procedures were carried out in accordance with the Declaration of Helsinki and according to respective institution guidelines³⁰. There were two data cohorts for the current investigation:

Training dataset -included a set of 13,580 clock drawings from participants age ≥ 65 years, primary English speaking, who completed clock drawing to command and copy conditions as part of routine medical care assessment in a preoperative setting.³¹ Data were collected from January 2018 to December 2019. Exclusion criteria were as follows: non-fluent in the English language; education < 4 years; visual, hearing, or motor extremity limitation that potentially inhibits the production of a valid clock drawing.

Classification dataset– consists of “*fine tuning*” and “*validation*” datasets with a set of individuals meeting criteria for dementia and a separate set of data from non-dementia peers. Dementia clocks are from individuals who had been evaluated as part of a community memory assessment program within Rowan University between February 2016 and March 2019. Individuals were seen by a neuropsychologist, a psychiatrist, and a social worker. Inclusion criteria: age ≥ 55 . Exclusion criteria: head trauma, heart

disease, or other major medical illness that can induce encephalopathy; major psychiatric disorders; documented learning disability; seizure disorder or other major neurological disorder; less than 6th-grade education, and history of substance abuse. All dementia participants were assessed with the Mini-Mental State Exam (MMSE), serum studies, and an MRI scan of the brain. Individuals have been described in prior research studies.³² As described in prior scientific papers, these individuals were diagnosed with either AD or VaD using standard diagnostic criteria, respectively.^{33,34}

Non-dementia peers had completed a research protocol of neuropsychological measures and neuroimaging with all data reviewed by two neuropsychologists. Inclusion criteria included age ≥ 60 , English as the primary language, and availability of intact activities of daily living (ADLs) as per Lawton & Brody's Activity of Daily Living Scale, completed by both the participant and their caregiver.³⁵ Exclusion criteria: clinical evidence of major neurocognitive disorder at baseline, as per the Diagnostic and Statistical Manual of Mental Disorders – Fifth Edition,³⁶ presence of a significant chronic medical condition, major psychiatric disorder, history of head trauma/neurodegenerative disease, documented learning disorder, epilepsy or other significant neurological illness, less than a 6th grade education, substance abuse in the past year, major cardiac disease, and chronic medical illness-induced encephalopathy. Participants were screened for dementia over the telephone using the Telephone Interview for Cognitive Status (TICS³⁷); and during an in-person interview with a neuropsychologist and a trained research coordinator who also evaluated comorbidity rating³⁸, anxiety, depression, ADLs, neuropsychological functioning, and digital clock drawing³⁹. Data from these participants were collected from September 2012 to November 2019. These data have been described elsewhere.^{2,18}

Procedure

Cohort participants completed two clock drawings – one to a command condition and another to a copy condition¹. The command condition required individuals to “Draw the face of a clock, put in all the numbers, and set the hands to ten after eleven”. The copy condition required individuals to draw the clock underneath a presented model of a clock. Drawings were completed using digital pen from Anoto, Inc. and associated smart paper^{16,17}. The digital pen technology captures and measures pen positioning on smart paper 75 times/second. 8.5 × 11inch smart paper folded in half giving participants a drawing area of 8.5 × 5.5inch. For the current project, only the final drawing was extracted and used for analyses.

Data from a training cohort of 13,580 unlabeled clock drawings to command and copy condition were used to train the VAE in an unsupervised manner. Thereafter, the trained VAE encoder network, which compresses clock drawings into a low-dimensional latent space, is fine-tuned for distinguishing dementia from control clock drawings. Command and copy clocks were not separated at this stage as we wanted the model to extract features from clock drawings in a way that is agnostic to any cognitive outcome. This serves to keep the extracted features general and useful for any downstream classification task. The latent features extracted by the VAE were passed to a classification network for distinguishing dementia

from control. This network was fine-tuned with 53 dementia and 60 control clock drawings and tested on 18 dementia and 20 control clock drawings. The fine-tuning and validation datasets were created by randomly shuffling the classification dataset and splitting it in a 3:1 ratio. Therefore, there was no demographic difference between the fine-tuning and validation dataset.

Individual clock sketches were extracted from file using contour detection. The extracted images were then flattened into one-dimensional vectors. These vectors were filtered to retain the ones with a size of less than 40,000 pixels (200 X 200). This filtering step is necessary to remove clock images containing excessive white space due to information sparsity issues. As clock drawings become sparser, the VAE tries to encode the interior white space of the clock instead of its drawn features, i.e., digits, hands, or circumference. This leads to poor latent space distributions that resemble white noise (Supplementary Fig S1), high reconstruction errors, and ineffective VAE encoder weights. Therefore, we restricted clock sizes to a maximum of 200 X 200 pixels. The filtered clocks were then resized to a fixed size of 100 X 100 and converted into 1-dimensional vectors (10000 X 1). These one-dimensional representations were used as inputs to the VAE. The preprocessing steps are illustrated in Supplementary Fig S2.

Models and Experimental setup

A variational autoencoder (VAE) is an unsupervised generative model with an encoding phase that projects input data onto a lower-dimensional latent space and a decoding phase that reconstructs the input data from random samples drawn from this latent space (Fig. 4). In the VAE model, the latent space distribution is created under the restriction that it follows a Gaussian distribution $N(Z_m, Z_s)$. This makes the VAE a generative model, as it can randomly sample this latent space distribution to create images resembling the input data which are not necessarily present in the input dataset. The use of a normal distribution as a prior does not lead to a loss of generality because the non-linear decoder network can mimic arbitrary input data distributions. No information about the generative features of clock drawing (e.g., total stroke length, clockface symmetry, coordinates of hands, number of digits, etc.) were supplied to the VAE network and we had no a priori expectations of making the latent dimension represent these generative features. We used unidimensional representations of clock drawings to train the VAE with an input dimension of 10,000, an intermediate dimension of 512, and an embedding dimension of 2. The embedding dimension was intentionally kept extremely low to inquire if such a low-dimensional manifold can extract meaningful clock features useful for classification. The model was trained for 50 epochs with a batch size of 16. The reconstruction loss of the VAE was chosen to be the binary cross-entropy loss. The trained latent representation of the VAE was used as input to a feed-forward fully connected neural network with a single hidden layer of 512 neurons for classification studies. Figure 4 shows the architecture of our networks and a conceptual workflow of our method. The top portion of the figure shows the training of the VAE. The bottom portion of the figure shows how the compressed latent space of the VAE, in the form of trained encoder weights, is used to create a task-specific classifier. The classifier network has a fully connected feed-forward neural network architecture. We fine-tuned the weights of this classifier using a small, annotated fine-tuning dataset to improve its

performance. The number of neurons in each layer of this classifier were finalized by using randomized grid search inside a 3-fold cross validation setting. The reported number of neurons gave the best average performance on the fine-tuning data over 3-fold cross-validation. Finally, the performance of this trained classifier was validated on the validation data and several important performance metrics namely, AUROC, Accuracy, Sensitivity, Specificity, Precision and Negative Predictive Value (NPV) were reported. The validation data was bootstrapped 100 times using random sampling with replacement to create confidence intervals of these metrics. The reported values constitute the median score, 2.5th quartile and 97.5th quartile of these metrics over the bootstrapped validation dataset.

Declarations

Acknowledgement and Funding

We would like to acknowledge Shawna Amini for managing the datasets used for this project.

This work was conducted at the University of Florida. C.P was supported by R01 NR014810 awarded by the National Institute of Nursing Research. C.P was also supported by R01 NS082386 awarded by the National Institute of Neurological Disorders and Stroke. C.P was supported by K07AG066813 by the National Institutes of Health. C.P and P.T were both supported by R01AG055337 by the National Institute on Aging, the National Center for Advancing Translational Science, and the University of Florida. P.T was supported by K07AG073468 by the National Institutes of Health. P.R was supported by National Science Foundation CAREER award 1750192.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute on Aging, National Institute of Nursing Research, National Institute of Neurological Disorders and Stroke, National Institutes of Health, National Center for Advancing Translational Science, or University of Florida.

Author contributions

C.P, P.T, D.L conceptualized the study. P.R, S.B designed the study. S.B acquired and analyzed the data. C.P, P.T, D.L, P.R, S.B and C.D interpreted the data. S.B and C.D drafted the manuscript. P.R, C.P, D.L substantively revised the manuscript. All authors approved the final version of the manuscript for submission. All authors have agreed both to be personally accountable for the author's own contributions and to ensure that questions related to the accuracy or integrity of any part of the work, even ones in which the author was not personally involved, are appropriately investigated, resolved, and the resolution documented in the literature.

Competing interests

The author(s) declare no competing interests.

Data Availability Statement

Datasets are available upon reasonable request. Reasonable requests will be reviewed to monitor compliance with the concerned authorities- National Institute of Health (NIH) and the Institutional Review Board (IRB). Relevant clinical trial numbers for the studies from which the datasets in this study have been constructed are NCT01986577 and NCT03175302.

References

1. Libon, D. J., Malamut, B. L., Swenson, R., Sands, L. P. & Cloud, B. S. Further analyses of clock drawings among demented and nondemented older subjects. *Arch Clin Neuropsychol* **11**, 193-205 (1996).
2. Dion, C. *et al.* Cognitive Correlates of Digital Clock Drawing Metrics in Older Adults with and without Mild Cognitive Impairment. *J Alzheimers Dis* **75**, 73-83, doi:10.3233/JAD-191089 (2020).
3. Freedman, M., Leach, L., Kaplan, E., Shulman, K. & Delis, D. C. *Clock drawing: A neuropsychological analysis.* (Oxford University Press, USA, 1994).
4. Cosentino, S., Jefferson, A., Chute, D. L., Kaplan, E. & Libon, D. J. Clock drawing errors in dementia: neuropsychological and neuroanatomical considerations. *Cogn Behav Neurol* **17**, 74-84, doi:10.1097/01.wnn.0000119564.08162.46 (2004).
5. Penney, D. *et al.* in *Annual Meeting of The International Neuropsychological Society.*
6. Piers, R. J. *et al.* Age and Graphomotor Decision Making Assessed with the Digital Clock Drawing Test: The Framingham Heart Study. *J Alzheimers Dis* **60**, 1611-1620, doi:10.3233/jad-170444 (2017).
7. Royall, D. R., Cordes, J. A. & Polk, M. CLOX: an executive clock drawing task. *J Neurol Neurosurg Psychiatry* **64**, 588-594, doi:10.1136/jnnp.64.5.588 (1998).
8. Shulman, K. I., Shedletsky, R. & Silver, I. L. The challenge of time: clock-drawing and cognitive function in the elderly. *International journal of geriatric psychiatry* **1**, 135-140 (1986).
9. Rouleau, I., Salmon, D. P., Butters, N., Kennedy, C. & McGuire, K. Quantitative and qualitative analyses of clock drawings in Alzheimer's and Huntington's disease. *Brain Cogn* **18**, 70-87, doi:10.1016/0278-2626(92)90112-y (1992).
10. Sunderland, T. *et al.* Clock drawing in Alzheimer's disease. A novel measure of dementia severity. *J Am Geriatr Soc* **37**, 725-729, doi:10.1111/j.1532-5415.1989.tb02233.x (1989).
11. Agrell, B. & Dehlin, O. The clock-drawing test. *Age and ageing* **27**, 399-404 (1998).
12. Shulman, K. I. Clock-drawing: is it the ideal cognitive screening test? *Int J Geriatr Psychiatry* **15**, 548-561, doi:10.1002/1099-1166(200006)15:6<548::aid-gps242>3.0.co;2-u (2000).
13. Spenciere, B., Alves, H. & Charchat-Fichman, H. Scoring systems for the Clock Drawing Test: A historical review. *Dement Neuropsychol* **11**, 6-14, doi:10.1590/1980-57642016dn11-010003 (2017).
14. Price, C. C. *et al.* Clock drawing in the Montreal Cognitive Assessment: recommendations for dementia assessment. *Dement Geriatr Cogn Disord* **31**, 179-187, doi:10.1159/000324639 (2011).
15. Frei, B. W. *et al.* Considerations for Clock Drawing Scoring Systems in Perioperative Anesthesia Settings. *Anesth Analg* **128**, e61-e64, doi:10.1213/ANE.0000000000004105 (2019).

16. Davis, R., Libon, D. J., Au, R., Pitman, D. & Penney, D. L. THink: Inferring Cognitive Status from Subtle Behaviors. *Proc Conf AAAI Artif Intell* **2014**, 2898-2905 (2014).
17. Souillard-Mandar, W. *et al.* Learning Classification Models of Cognitive Conditions from Subtle Behaviors in the Digital Clock Drawing Test. *Mach Learn* **102**, 393-441, doi:10.1007/s10994-015-5529-5 (2016).
18. Davoudi, A. *et al.* Classifying Non-Dementia and Alzheimer's Disease/Vascular Dementia Patients Using Kinematic, Time-Based, and Visuospatial Parameters: The Digital Clock Drawing Test. *J Alzheimers Dis* **82**, 47-57, doi:10.3233/JAD-201129 (2021).
19. Binaco, R. *et al.* Machine Learning Analysis of Digital Clock Drawing Test Performance for Differential Classification of Mild Cognitive Impairment Subtypes Versus Alzheimer's Disease. *J Int Neuropsychol Soc* **26**, 690-700, doi:10.1017/S1355617720000144 (2020).
20. Souillard-Mandar, W. *et al.* DCTclock: Clinically-Interpretable and Automated Artificial Intelligence Analysis of Drawing Behavior for Capturing Cognition. *Front Digit Health* **3**, 750661, doi:10.3389/fdgth.2021.750661 (2021).
21. Gomes-Osman, J. *et al.* Aging in the Digital Age: Using Technology to Increase the Reach of the Clinician Expert and Close the Gap Between Health Span and Life Span. *Frontiers in Digital Health* **3** (2021).
22. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
23. Kingma, D. & Welling, M. (2019).
24. Kingma, D. P., Mohamed, S., Rezende, D. J. & Welling, M. in *Advances in neural information processing systems*. 3581-3589.
25. Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K. & Winther, O. How to train deep variational autoencoders and probabilistic ladder networks. *arXiv preprint arXiv:1602.02282* **3** (2016).
26. Kim, H. & Mnih, A. in *International Conference on Machine Learning*. 2649-2658 (PMLR).
27. Kim, M., Wang, Y., Sahu, P. & Pavlovic, V. Relevance factor VAE: Learning and identifying disentangled factors. *arXiv preprint arXiv:1902.01568* (2019).
28. Kim, M., Wang, Y., Sahu, P. & Pavlovic, V. in *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2979-2987.
29. Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
30. Moons, K. G. *et al.* Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis (TRIPOD): explanation and elaboration. *Ann Intern Med* **162**, W1-73, doi:10.7326/M14-0698 (2015).
31. Amini, S. *et al.* Feasibility and Rationale for Incorporating Frailty and Cognitive Screening Protocols in a Preoperative Anesthesia Clinic. *Anesth Analg* **129**, 830-838, doi:10.1213/ANE.0000000000004190 (2019).

32. Emrani, S. *et al.* Alzheimer's/vascular spectrum dementia: classification in addition to diagnosis. *Journal of Alzheimer's Disease* **73**, 63-71 (2020).
33. Price, C. C., Jefferson, A. L., Merino, J. G., Heilman, K. M. & Libon, D. J. Subcortical vascular dementia: integrating neuropsychological and neuroradiologic data. *Neurology* **65**, 376-382, doi:10.1212/01.wnl.0000168877.06011.15 (2005).
34. Price, C. C. *et al.* Leukoaraiosis severity and list-learning in dementia. *Clin Neuropsychol* **23**, 944-961, doi:10.1080/13854040802681664 (2009).
35. Lawton, M. P. & Brody, E. M. Assessment of older people: self-maintaining and instrumental activities of daily living. *Gerontologist* **9**, 179-186 (1969).
36. American Psychiatric Association, D. & Association, A. P. (Washington, DC: American Psychiatric Association, 2013).
37. Welsh, K. A., Breitner, J. C. & Magruder-Habib, K. M. Detection of dementia in the elderly using telephone screening of cognitive status. *Neuropsychiatry, Neuropsychology, & Behavioral Neurology* (1993).
38. Charlson, M. E., Pompei, P., Ales, K. L. & MacKenzie, C. R. A new method of classifying prognostic comorbidity in longitudinal studies: development and validation. *J Chronic Dis* **40**, 373-383, doi:10.1016/0021-9681(87)90171-8 (1987).
39. Davis, R. *et al.* The Digital Clock Drawing Test (dCDT) I: Development of a new computerized quantitative system. *The International Neuropsychological Society* (2011).

Figures

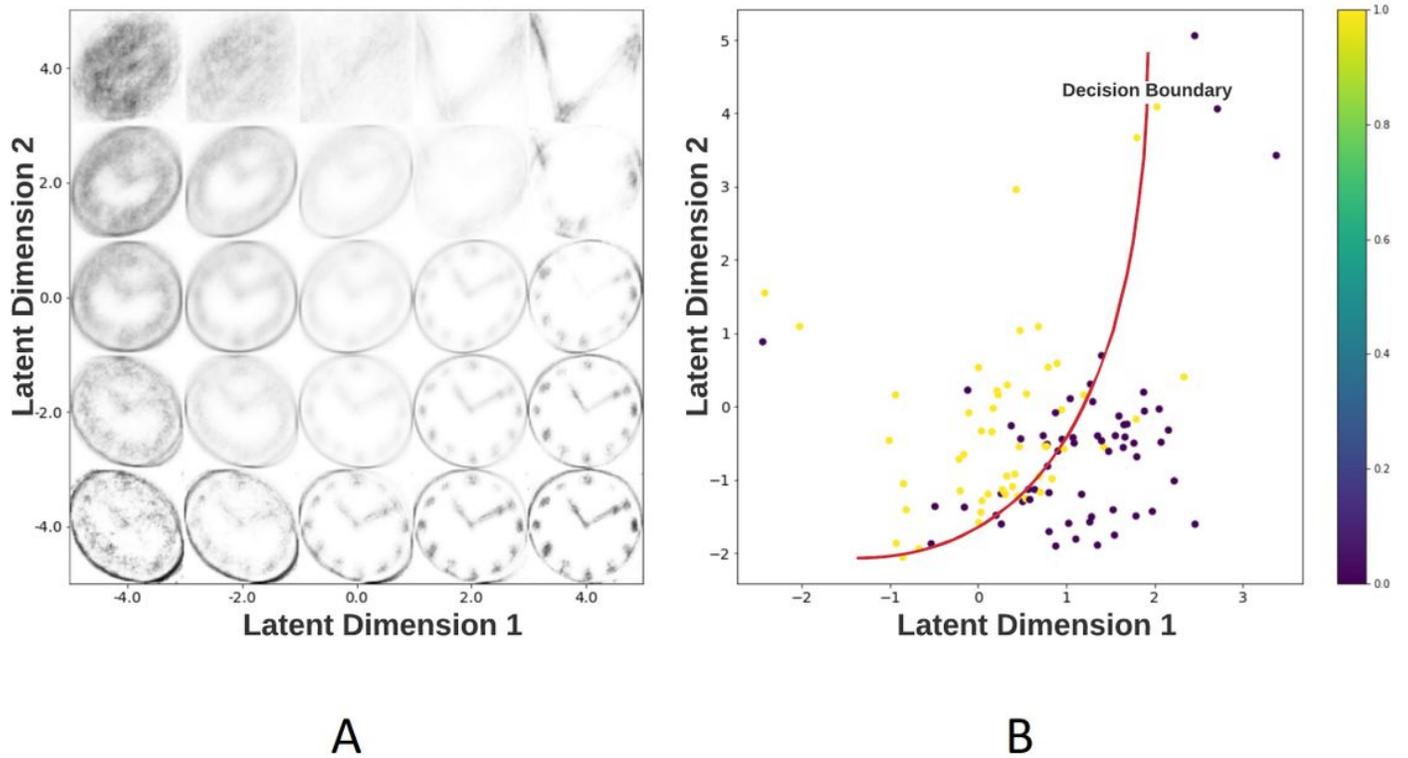


Figure 1

Distribution of reconstructed clocks in the VAE latent space. A. Clock drawing reconstructions are represented as a function of the two VAE latent dimensions. This shows the variety of reconstructions generated by the VAE to capture the distribution of the training dataset. **B.** Scatterplot showing the distribution of the latent vectors belonging to clocks in the fine-tuning dataset divided into dementia (=1) and control (=0) groups. The red curve represents a possible decision boundary between the two groups. Using our neural network classifier, we aim to learn such a decision boundary to classify the two groups.

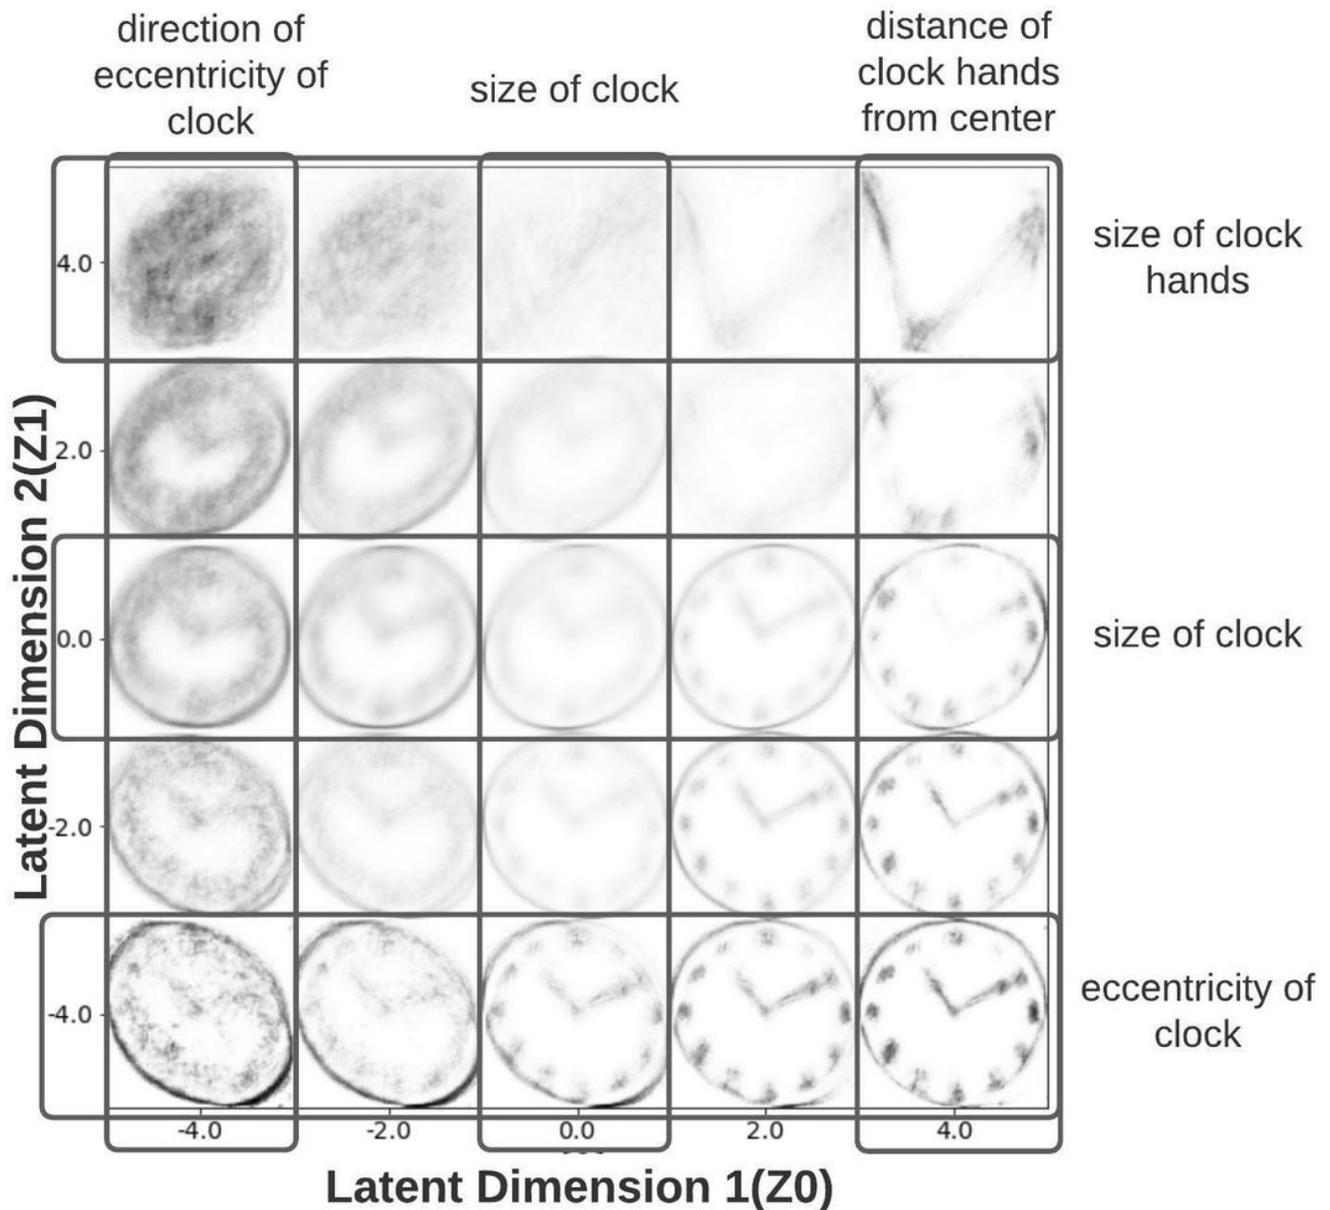


Figure 2

Generative factors of clock drawing are discovered using traversals over VAE latent space. Different regions of the two-dimensional manifold constructed by the VAE map to specific artistic features/anomalies of clock drawings. Latent Dimension 1=Z0, Latent Dimension 2=Z1. Left side of the VAE latent space: direction of eccentricity of reconstructed clocks reverses from left to right as Z1 increases from -4 to +4 given Z0 = -4. This change is correlated with a decrease in clock size shown by darkening of the reconstructed clock. Right side of the VAE latent space: distance of intersection point of clock hands from the geometrical clock center increases as Z1 changes from -4 to +4 given Z0 = 4. This change is correlated with a loss of the circular clock boundary. Top half of the VAE latent space: size of the clock hands increases as Z0 changes from -4 to +4 given Z1 = 4. This change is correlated with the

loss of clock face boundary. Bottom half of the VAE latent space: angle of eccentricity of the clock decreases as Z_0 changes from -4 to +4 with $Z_1 = -4$. This change is correlated with an increase in clock size. X-axis of the VAE latent space: size of the clock increases as Z_0 changes from -4 to +4 with $Z_1 = 0$.

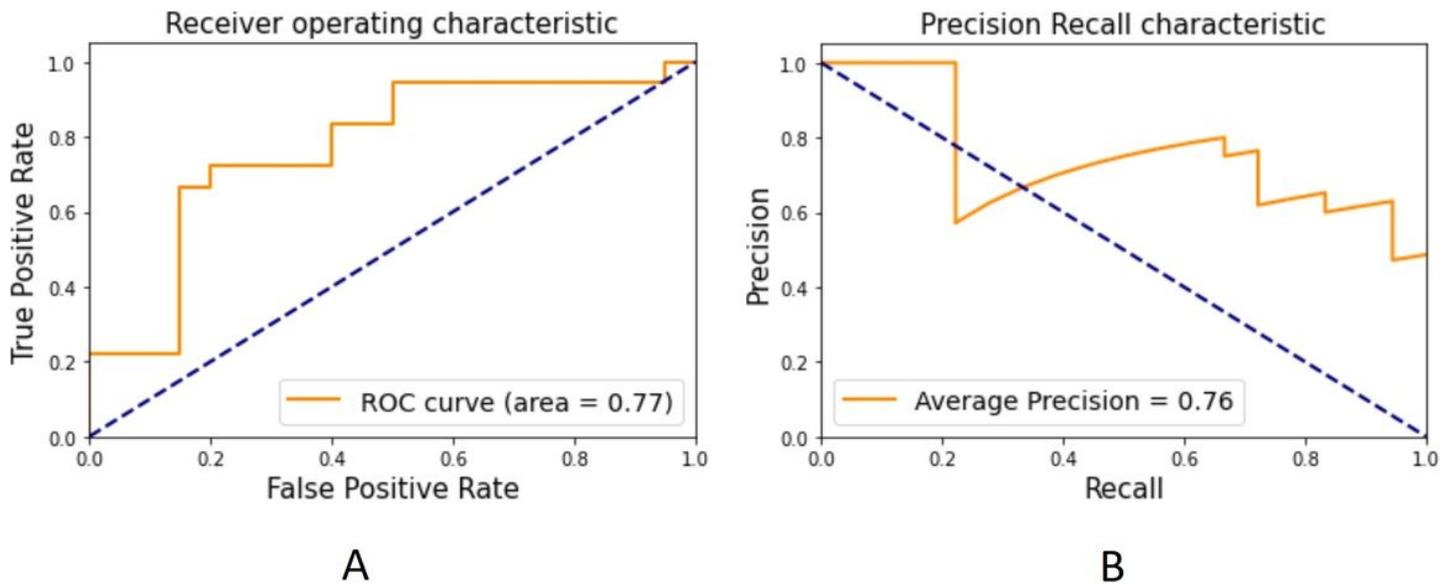


Figure 3

Performance of the classifier on test dataset. **A.** The test dataset consists of 18 dementia and 20 control subjects. Area under Receiver Operating characteristic is 0.77. **B.** Area under Precision-Recall curve is 0.76.

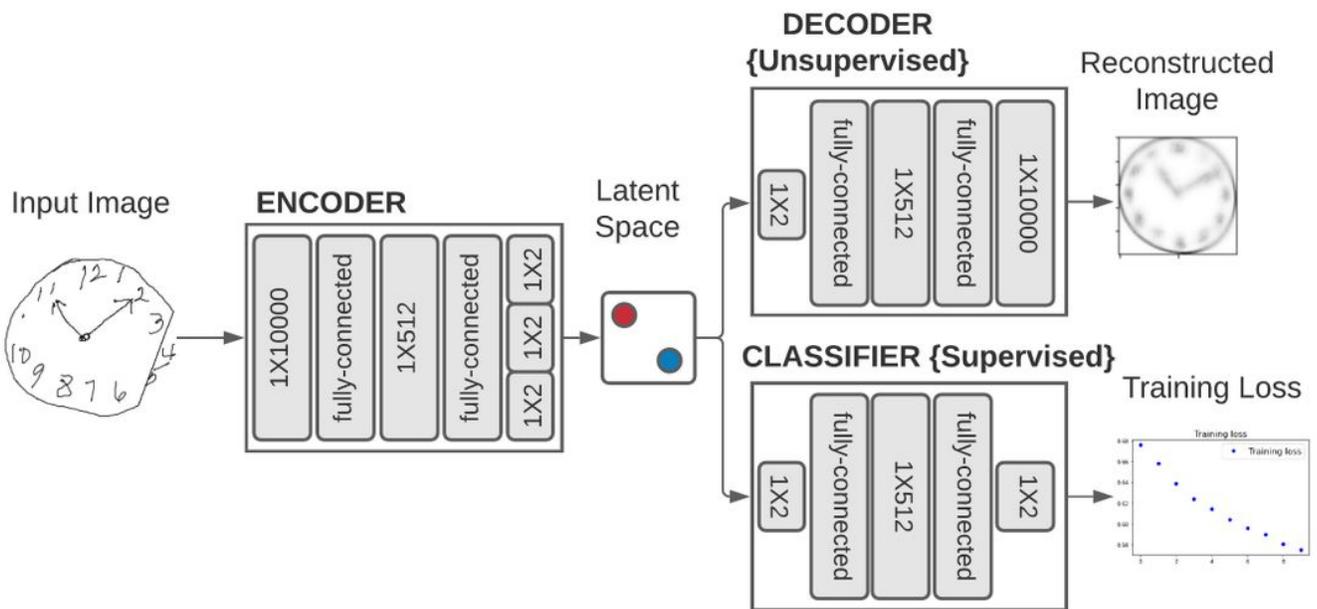


Figure 4

Conceptual workflow for training VAE with unlabeled dataset and constructing classifiers with latent space of VAE. Clock images are passed into the VAE encoder in the form of a 1X10000 vector. The top part of the schema represents the VAE decoder-encoder couplet which is trained to minimize the reconstruction loss of the clock drawing using an information bottleneck of a two-dimensional latent space. The lower part of the figure shows how the latent dimensions which capture a compressed representation of a clock drawing are passed into a classifier that mimics the architecture of the VAE decoder. This classifier is fine-tuned to reduce the loss in predicting dementia.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [NatureScReportsSupplementaryInfo.docx](#)
- [SupplementaryFigS1.png](#)
- [SupplementaryFigS2.png](#)
- [SupplementaryVideoV1.gif](#)
- [SupplementaryVideoV2.gif](#)
- [SupplementaryVideoV3.gif](#)
- [SupplementaryVideoV4.gif](#)
- [SupplementaryVideoV5.gif](#)