The AI Neuropsychologist: Automatic scoring of memory deficits with deep learning

AUTHORS:

Nicolas Langer[1,2,3*], Maurice Weber[4], Bruno Hebling Vieira[1,2,3], Dawid Strzelczyk[1,2,3], Lukas Wolf[1], Andreas Pedroni[1,2,3], Jonathan Heitz[4], Stephan Müller[4], Christoph Schultheiss[4], Marius Tröndle[1,2,3], Juan Carlos Arango Lasprilla[5,6], Diego Rivera[7,8], Federica Scarpina[9,10], Qianhua Zhao[11], Rico Leuthold[12], Flavia Wehrle[13], Oskar G. Jenni[13], Peter Brugger[14], Tino Zaehle[15], Romy Lorenz[16,17,18], Ce Zhang[4]

AFFILIATIONS:
1: Methods of Plasticity Research, Department of Psychology, University of Zurich, Zurich, Switzerland
2: University Research Priority Program (URPP) Dynamics of Healthy Aging, Zurich, Switzerland
3: Neuroscience Center Zurich (ZNZ), University of Zurich and ETH Zurich, Zurich, Switzerland
4. Department of Computer Science, ETH, Zurich, Switzerland
5. Giunti Psychometrics. Florence, Italy.
6. Virginia Commonwealth University. Richmond, Virginia.
7. Department of Health Science, Public University of Navarre, Pamplona, Spain
8. Instituto de Investigación Sanitaria de Navarra (IdiSNA), Pamplona, Spain
9. "Rita Levi Montalcini" Department of Neurosciences, University of Turin, Italy
10. I.R.C.C.S. Istituto Auxologico Italiano, U.O. di Neurologia e Neuroriabilitazione, Ospedale San Giuseppe, Piancavallo (VCO), Italy
11. Huashan Hospital, Shanghai, China
12. Smartcode, Zurich, Switzerland
13. University Children's Hospital Zurich, Child Development Center, Zurich, Switzerland
14. Rehabilitation Center, Valens, Switzerland
15. Universitätsklinik für Neurologie, Universitätsklinikum Magdeburg, Magdeburg, Germany
16. MRC Cognition and Brain Sciences Unit, University of Cambridge, Cambridge, United Kingdom
17. Department of Neurophysics, Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany
18. Stanford University, Stanford, CA

CORRESPONDING AUTHOR(*):

Nicolas Langer
Methods of Plasticity Research
Department of Psychology
University of Zurich
Andreasstrasse 15
CH-8050, Zürich, Switzerland
E-mail: n.langer@psychologie.uzh.ch
Tel.: +41 44 635 34 14

Abstract

Memory deficits are a hallmark of many different neurological and psychiatric conditions. The Rey-Osterrieth complex figure (ROCF) is the state–of-the-art assessment tool for neuropsychologists across the globe to assess the degree of non-verbal visual memory deterioration. To obtain a score, a trained clinician inspects a patient's ROCF drawing and quantifies deviations from the original figure. This manual procedure is time-consuming, slow and scores vary depending on the clinician's experience, motivation and tiredness. Here, we leverage novel deep learning architectures to automatize the rating of memory deficits. For this, a multi-head convolutional neural network was trained on 20225 ROCF drawings. Unbiased ground truth ROCF scores were obtained from crowdsourced human intelligence. The neural network outperforms both online raters and clinicians. Our AI-powered scoring system provides healthcare institutions worldwide with a digital tool to assess objectively, reliably and time-efficiently the performance in the ROCF test from hand-drawn images.

Introduction

Neurological and psychiatric disorders are among the most common and debilitating illnesses across the lifespan. In addition, the aging of our population, with the increasing prevalence of physical and cognitive disorders, poses a major burden on our society with an estimated economic cost of 2.5 trillion US$ per year (Trautmann, Rehm, and Wittchen 2016). Currently, neuropsychologists typically use paper-pencil tests to assess individual neuropsychological functions and brain dysfunctions, including memory, attention, reasoning and problem-solving. Most neuropsychologists around the world use the Rey-Osterrieth complex figure (ROCF) in their daily clinical practice (Rabin, Barr, and Burton 2005; Rabin, Paolillo, and Barr 2016), which provides insights into a person's nonverbal visuo-spatial memory capacity in healthy and clinical populations of all ages, from childhood to old age (Shin et al. 2006a).

Our estimation revealed that a single neuropsychological division (e.g. at the University Hospital Zurich) scores up to 6000 ROCF drawing per year. The ROCF test has several advantages as it does not depend on auditory processing and differences in language skills that are omnipresent in cosmopolitan societies (Osterrieth 1944; Somerville, Tremont, and Stern 2000). The test exhibits adequate psychometric properties (e.g. sufficient test–retest reliability (John E. Meyers and Volbrecht 1999; Levine et al. 2004) and internal consistency (Berry, Allen, and Schmitt 1991; Fastenau, Bennett, and Denburg 1996)). Furthermore, the ROCF test has demonstrated to be sensitive to discriminate between various clinical populations (Alladi et al. 2006) and the progression of Alzheimer's disease (Trojano and Gainotti 2016).

The ROCF test consists of three test conditions: First, in the *Copy condition* subjects are presented with the ROCF and are asked to draw a copy of the same figure. Subsequently, the ROCF figure and the drawn copy are removed and the subject is immediately instructed to reproduce the figure from memory (*Immediate Recall condition*). After a delay of 30 minutes, the subject is required to draw the same figure once again (*Delayed Recall* condition). For further description of the clinical usefulness of the ROCF please refer to (Shin et al. 2006a).

The current quantitative scoring system (J. E. Meyers, Bayless, and Meyers 1996) splits the ROCF into 18 identifiable elements (see Figure 1A), each of which is considered separately and marked on the accuracy in both distortion and placement according to a set of rules and scored between 0 and 2 points (see supplementary Figure 3). Thus, the scoring is currently undertaken manually by a trained clinician, who inspects the reproduced ROCF drawing and tracks deviations from the original figure, which can take up to 15 minutes per figure (individually *Copy*, *Immediate Recall* and *Delayed Recall* conditions).

One major limitation of this quantitative scoring system is that the criteria of what position and distortion is considered “accurate” or “inaccurate” may vary from clinician to clinician (Groth-Marnat 2000; Shin et al. 2006b; Canham, Smith, and Tyrrell 2000a). In addition, the scoring might vary as a function of motivation and tiredness or because the clinicians may be unwittingly influenced by interaction with the patient. Therefore, an automated system that offers reliable, objective, robust and standardized scoring, while saving clinicians’ time, would be desirable from an economic perspective and more importantly leads to more accurate scoring and subsequently diagnosing.

In recent years, computerized image analysis and machine learning methods have entered the clinical neuropsychological diagnosis field providing the potential for establishing quantitative scoring methods. Using machine-learning methods, such as convolutional neural networks and support vector machines, studies have successfully recognized visual structures of interest produced by subjects in the Bender Gestalt Test (BGT) (Nazar et al. 2017), and the Clock Draw Task (CDT), (Kim, Cho, and Do 2010; Harbi, Hicks, and Setchi 2016) - both are less frequently applied screening tests for visuospatial and visuo-constructive (dis-)abilities (Webb et al. 2021). Given the wide application of the ROCF, it is not surprising that we are not the first to take steps towards a machine-based scoring system. (Canham, Smith, and Tyrrell 2000b) have developed an algorithm to automatically identify a selection of parts of the ROCF (6 of 18 elements) with great precision. This approach provides first evidence for the feasibility of automated segmentation and feature recognition of the ROCF. More recently, (Vogt et al. 2019) developed an automated scoring using a deep neural network. The authors reported a .88 Pearson correlation with human ratings, but equivalence testing demonstrated that the automated scoring did not produce strictly equivalent total scores compared to the human ratings. Moreover, it is unclear if the reported correlation was observed in an independent test data set, or in the training set. Finally, (Petilli et al. 2021) have proposed a novel tablet-based digital system for the assessment of the ROCF task, which provides the opportunity to extract a variety of parameters such as spatial, procedural, and kinematic scores. However, none of the studies described have been able to produce machine-based scores according to the original scoring system currently used in clinics (J. E. Meyers, Bayless, and Meyers 1996) that are equivalent or superior to human raters. A major challenge in developing automated scoring systems is to gather a rich enough set of data with instances of all possible mistakes that humans can make. In this study, we have trained convolutional neural networks with over 20’000 digitized ROCFs from different populations regarding age and diagnostic status (healthy individuals or individuals with neurological and psychiatric disorder (e.g. Alzheimer, Parkinson)).

Results

Human Rating Performance: A human MSE

For our experiments, we used a data set of 20’225 hand-drawn ROCF images collected from research and clinical settings in various countries spanning the complete range of scores (Figure 1). Our analysis (see section: Clinicians’ Scoring) and previous work (Canham, Smith, and Tyrrell 2000a) suggested that the scoring conducted by clinicians may not be consistent, because the clinicians may be unwittingly influenced by the interaction with the patient/participant or the scoring might vary as a function of their motivation and tiredness. For this reason, we have harnessed a large pool (~5000) of human workers (crowdsourced human intelligence). In a self-developed interactive web application, we first trained the human workers by guided instructions and examples. An average of 13.38 (sd = 2.23) crowdsource-based scores per figure were gathered. For each drawing, the ground truth score was obtained by computing the median scores for every item in the figure, which were then summed to yield the final total score of the image. Assuming that the sum of median scores of all items of an image is the ground truth, the average rating performance of human raters can be estimated by computing the mean squared error (MSE) and the mean absolute error (MAE) of the human ratings by first computing the mean error for each human rater, and then computing the mean of all individual MSEs or MAEs. The *average human MSE* over all images is 16.3, and the *average human MAE* is 2.41.

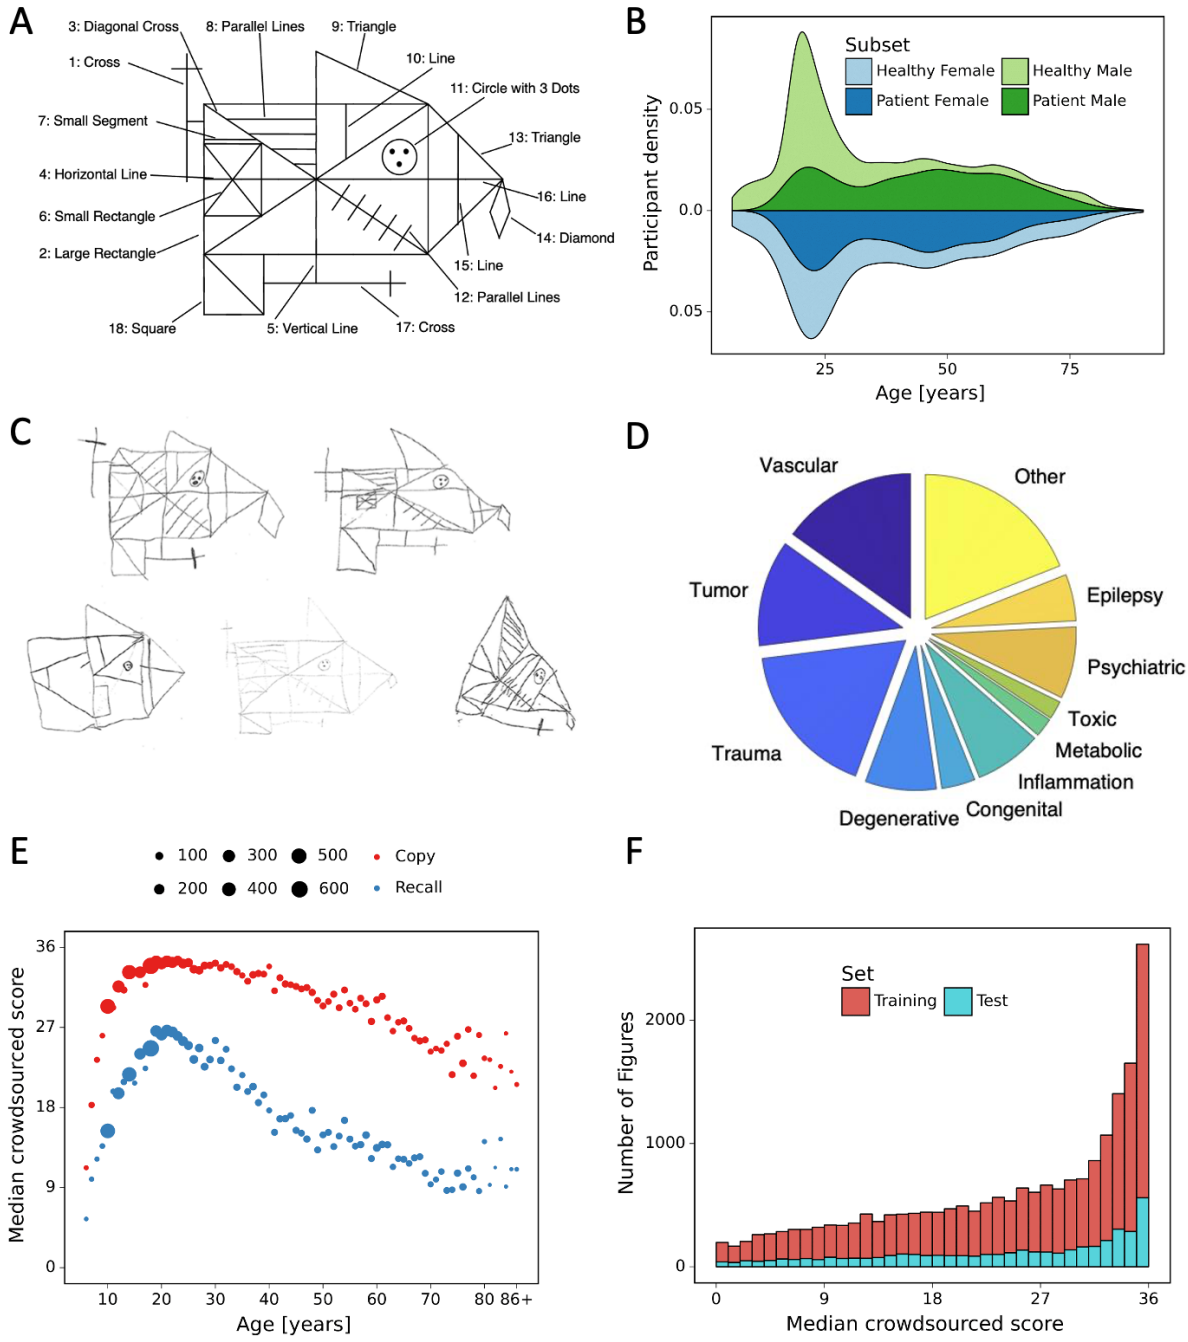

Figure 1. A: ROCF figure with 18 elements. B: demographics of the participants and clinical population. C: examples of hand-drawn ROCF images. D: Pie chart illustrates the numerical proportion of the different clinical conditions. E: performance in the copy and (immediate) recall condition across the lifespan in the present data set. F: distribution of number of images for each total score (online raters).

Clinicians' Scoring

For a subset (4030) of the 20225 ROCF images, we had access to scores conducted by different professional clinicians. This enabled us to analyze and compare the scoring of professionals to the crowdsourcing results. The clinician MSE over all images is 9.25 and the clinician MAE is 2.15, indicating a better performance of the clinicians compared to the average human rater. Additional analyses revealed that clinicians systematically slightly overestimated the scores compared to the ground truth (Figure 3E). For instance, the mean of clinicians' ratings is 22.4 whereas the mean of the ground truth labels (median from crowdsourcing) of the 4030 images clinicians rated is just 21.1. This was further tested by permutation statistics using 10'000 permutations, which yields this difference to be significant ($p<0.0001$). Even after we correct for outliers, by ignoring the images in which clinicians and online raters differ by more than ten points (i.e. 34 figures), the difference is still existent ($p<0.0001$).

Machine-based scoring system

For the automated scoring of ROCF drawings, two variations of deep learning system were implemented: a regression-based network and a multilabel classification network. The multilabel classification problem is based on the assumption that one drawing can contain multiple classes. That is, the scores for a specific element correspond to four mutually exclusive classes (i.e. 0, 0.5, 1, 2), which corresponds to the original Osterrieth scoring system (J. E. Meyers, Bayless, and Meyers 1996) (see also supplementary Figure 3). Each class can appear simultaneously with score classes corresponding to other elements in the figure. While the scores for different elements share common features such as e.g., texture or drawing style, which are relevant for the element score, one can view the scoring of different elements as independent processes, given these shared features. These considerations are the starting point for the architectural design of our scoring system, shown in Figure 2A. The model consists of a shared feature extractor, followed by 18 element classifiers, each taking as input the shared features extracted by the primary feature extractor. For a more detailed description of our architecture, we refer the reader to the methods section.

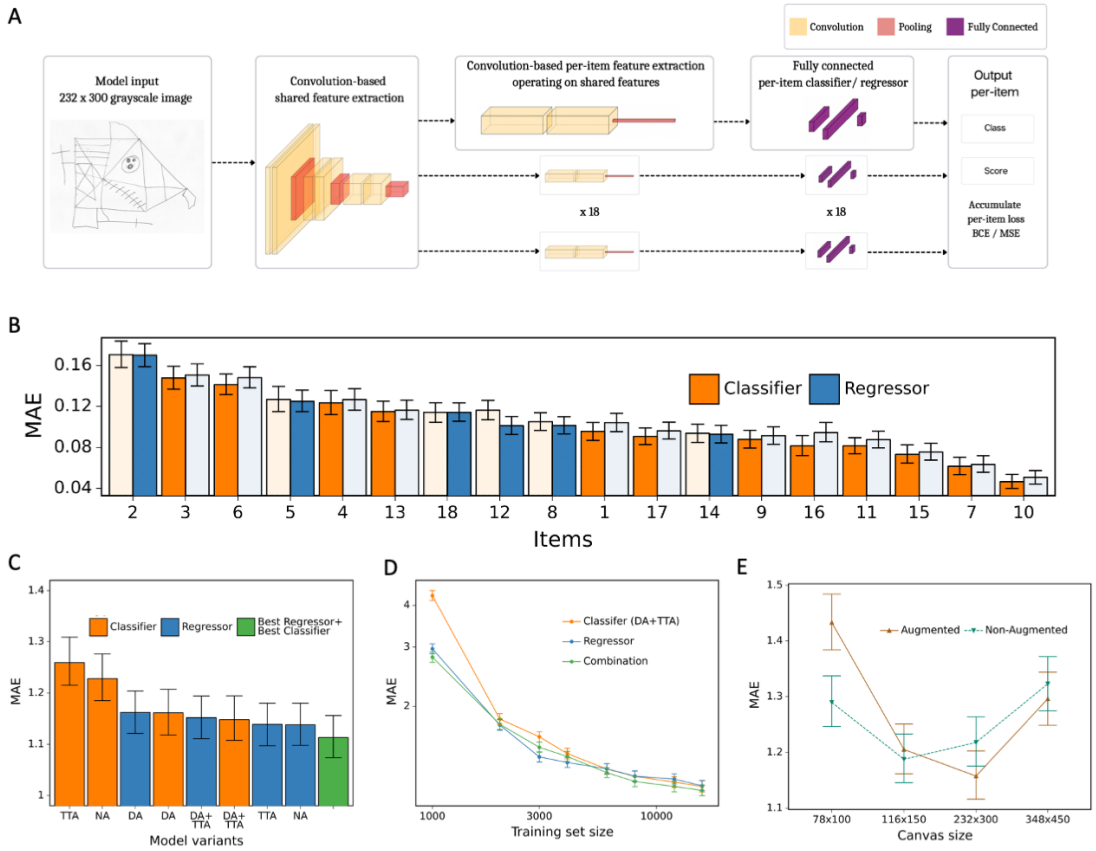

Figure 2. A: network architecture, constituted of a shared feature extractor and 18 item-specific feature extractors and output blocks. The shared feature extractor consists of three convolutional blocks, whereas item-specific feature extractors have one convolutional block with global max-pooling. Convolutional blocks consist of two convolution and batch-normalization pairs, followed by max-pooling. Output blocks consist of two fully connected layers. ReLU activation is applied after batch normalization. After pooling, dropout is applied. B. item-specific MAE for the regression-based network (blue) and multilabel classification network (orange). In the final model, we determine whether to use the regressor or classifier network based on its performance in the validation data set, indicated by an opaque color in the bar chart. In case of identical performance, the model resulting in least variance was selected. C: Model variants were compared and the performance of the best model (green) is displayed; Clf: multilabel classification network; Reg: regression-based network; NA: no augmentation; DA: data augmentation; TTA: test time augmentation. D. Convergence analysis revealed that after ~8000 images, no substantial improvements could be achieved by including more data (see also supplementary material). E. The effect of image size on the model performance measured in terms of MAE.

To train and evaluate our model, we used 80% of the dataset for training and validation (20% of the training data); and 20% as a testing set, which was never used for model learning. Before feeding the drawings to our model, each drawing in the dataset was converted to a grayscale image and resized to fit within a 232 x 300 pixels canvas, maintaining the aspect ratio to avoid distortions. In Figure 2E, we show the results of a study about the ideal image size, revealing that both larger and smaller image sizes do not bring any significant advantage to the model in terms of total score error. The scoring system trained with this setup achieved a total score MSE of 3.56 and MAE of 1.22, which is already considerably lower than the corresponding human performance (average human MSE is 16.3, MAE is 2.41).

In addition, to make our model more robust against geometric and semantic-preserving transformations like rotations and changes in perspective, we implemented a data augmentation pipeline consisting of such transformations. This has the additional benefit of enhancing our training set and is effective against overfitting to the training set. This has led to a further improvement in accuracy with an MSE of 3.37 and MAE of 1.16.

Furthermore, we investigated test time augmentation as it has been shown to improve classification accuracy in various realistic settings (Moshkov et al. 2020; Amiri et al. 2020; Krizhevsky, Sutskever, and Hinton 2017). As our results indicate, this can also be observed in our setting where combining data augmentation with test time augmentation leads to a further improvement of performance with an MSE of 3.32 and MAE of 1.15.

One can also view the task as being akin to a regression problem since there is a close relation between neighboring scores for a specific element. Indeed, a classification model is penalized equally when predicting score 0.0 as it is when predicting 1.0 for an element for which the true score is 2.0. However, the latter prediction is clearly “less wrong” from a practical point of view. To incorporate this information, we have converted the original multilabel classification model to a multilabel regression model, where each element regressor is trained with an MSE loss and outputs a single score. Analogous to the classification approach, the element regression models are built on top of a shared feature extractor. This has led to a further slight improvement in performance, reducing the total score MSE to 3.07 and the MAE to 1.14.

Finally, our best model results from combining the regression model with the multilabel classification model in the following way: For each item of the figure we determine whether to use the regressor or classifier based on its performance on the validation set (Figure 2B). Aggregating the two models in this way, leads to a MSE of 3.00 and a MAE of 1.11 (see supplementary table S1 for the performance metrics of all model variants and table S2 for the

item-specific MSE, MAE and R-squared). Figure 2C presents the error for all combinations of data and test time augmentation with the multilabel classification and regression model separately and in combination. During the experiments it became apparent that for the multilabel classification network applying both data augmentation as well as test-time-augmentation jointly improves the model's performance (see Figure 2C). Somewhat surprisingly, applying these elements to the multihead regression model did not improve the performance compared to the non-augmented version of the model.

We further conducted a more fine-grained analysis of our best model. Figure 3A shows the score for each figure in the data set contrasted with the ground truth score (i.e. median of online raters). In addition, we computed the difference between ground truth score and predicted score revealing that errors made by our model are concentrated closely around 0 (Figure 3B), while the distribution of clinician's errors is much more spread out (Figure 3D and 3E). It is interesting to note that, like the clinicians, our model exhibits a slight bias towards higher scores, although much less pronounced. Importantly, the model does not demonstrate any considerable bias towards specific figure elements. In contrast to the clinicians, the MAE is very balanced across each individual item of the figure (Figure 3C and 3F, and supplementary table S2). Finally, a detailed breakdown of the expected performance across the entire range of total scores is displayed for the model (Figure 3G and supplementary table S3), the clinicians (3H) and average online raters (3I). As can be seen, the model exhibits a comparable MAE for the entire range of total scores, although there is a trend that high scores exhibit lower MAEs.

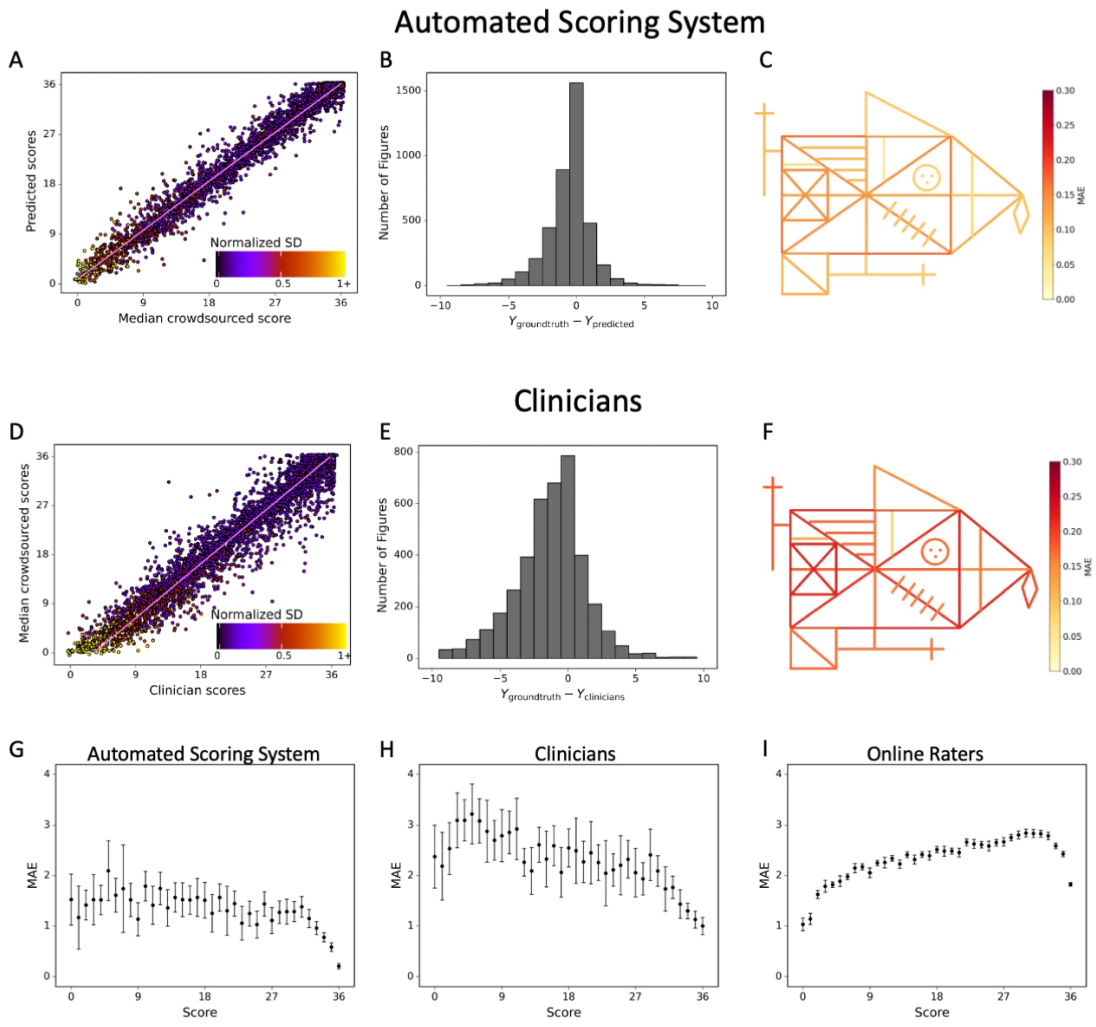

Figure 3. Contrasting the ratings of our model (A) and clinicians (D) against the ground truth. revealed a larger deviation from the regression line for the clinicians. A jitter is applied to better highlight the dot density. The distribution of errors for our model (B) and the clinicians ratings (E) are displayed. The MAE of our model (C) and the clinicians (F) is displayed for each individual item of the figure (see also supplementary table S2). The model performance across the entire range of total scores for model (G), clinicians (H) and online raters (I) are presented.

Robustness Analysis

Using data augmentation did not critically improve the accuracy, which is likely due to the fact that our dataset is already large and diverse enough. However, data augmentation does significantly improve robustness against semantic transformations, as can be seen from supplementary Figure S1. In particular, using our data augmentation pipeline, the model becomes much more robust against rotations and changes in perspective. On the other hand, for changes to brightness and contrast we could not find a clear trend. This is however not surprising as the data augmentation pipeline does not explicitly encourage the model to be more robust against these transformations. Overall, we demonstrate that our scoring system is highly robust for realistic changes in rotations, perspective, brightness, and contrast of the images (Figure 4). Performance is degraded only under unrealistic and extreme transformations.

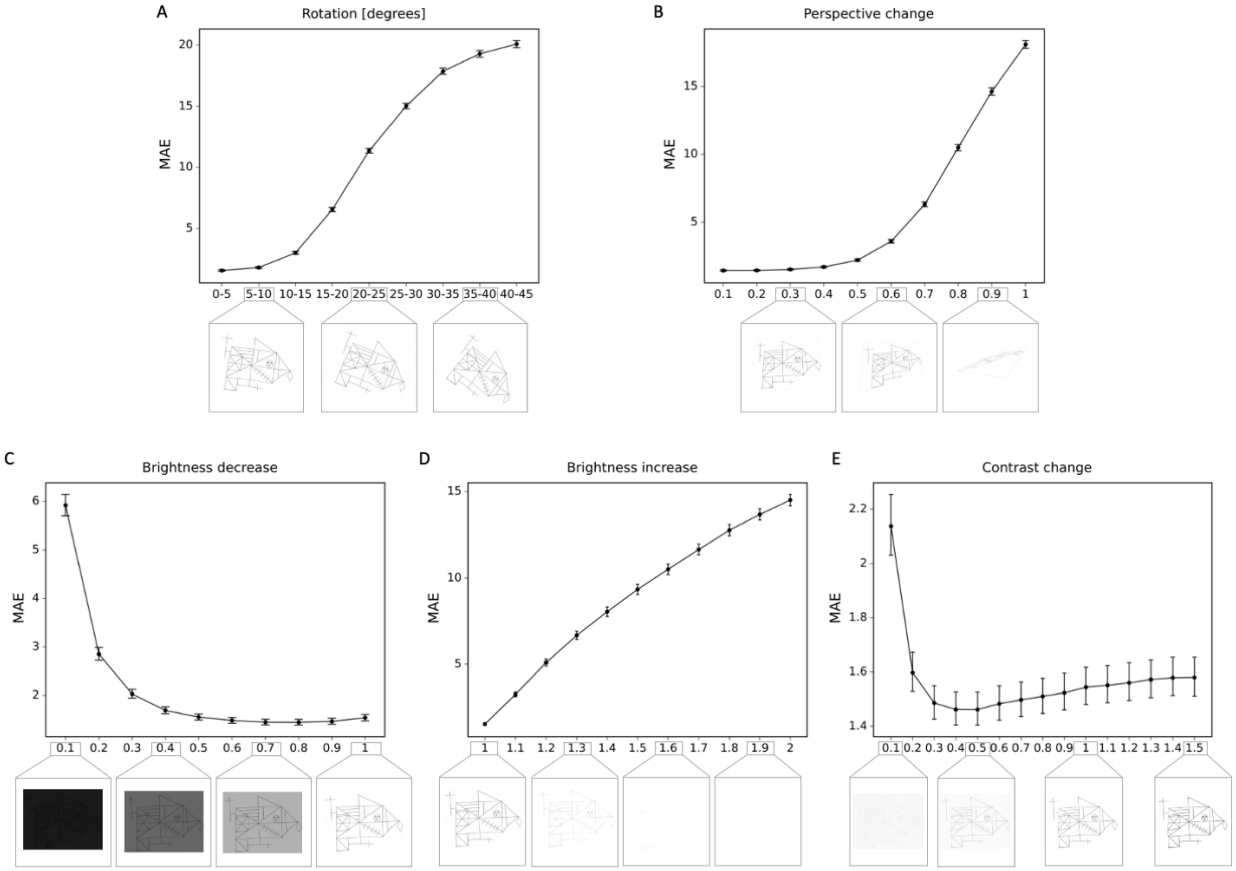

Figure 4. Robustness to geometric, brightness and contrast variations. The MAE is depicted for different degrees of transformations. In addition examples of the transformed ROCF draw are provided.

Discussion

In this study, we developed an AI-based scoring system for a nonverbal visuo-spatial memory test that is being utilized in clinics around the world on a daily basis. For this, we trained two variations of deep learning systems and used a rich data set of 20’225 hand-drawn ROCF images covering the entire life span that was collected from six different research and clinical environments. By leveraging human crowdsourcing we obtained unbiased high-precision training labels. Our best model results from combining a multihead convolutional neural network for regression with a network for multilabel classification. Data and test time augmentation were used to improve the accuracy and made our model more robust against geometric transformations like rotations and changes in perspective.

Overall, our results provide evidence that our AI-scoring tool outperforms both amateur raters and clinicians. Our model performed highly unbiased as it yielded predictions very close to the ground truth and the error was similarly distributed around zero. The scoring system reliably identified and accurately scored individual figure elements in previously unseen ROCF drawings, which facilitates explainability of the AI-scoring system. Comparable performance for each figure element indicates no considerable bias towards specific figure compositions. Furthermore, the error of our scoring system is rather balanced across the entire range of total scores. An exception are the highest scored images displaying the smallest error, which has two explanations: First, compared to low score images, for which the same score can be achieved by numerous combinations of figure elements, the high score images by nature do not have as many possible combinations. Secondly, the data set contains a proportional larger amount of available training data for the high score images.

While the ROCF is one of the most commonly employed neuropsychological test to evaluate nonverbal visuo-spatial memory in clinical setting (Shin et al. 2006c), the current manual scoring in clinics is time-consuming (5 to 15 minutes for each drawing), requires training, and ultimately depends on subjective judgements, which inevitably introduces human scoring biases (Watkins 2017; Franzen 2000). This has been criticized for its unreliability by various experts in the field (Tupler et al. 1995; Fastenau, Bennett, and Denburg 1996). While extensive training courses can improve inter-rater reliability (Franzen 2000; Huygelier et al. 2020) and previous studies (on very small sample size) reported human inter-rater reliability as high as 0.9 (J. E. Meyers,

Bayless, and Meyers 1996), our own analysis on a much larger data set (see supplementary Figure S4) demonstrates that the published inter-rater reliability is overestimated and potentially derived from clinicians, which are highly skilled and were trained in an identical fashion. This is probably rarely the case across different clinics, which precludes large-scale comparisons of ROCF test data, including multiple raters. Thus, an objective and reliable automated scoring system is in great demand as it can overcome the problem of intra- and inter-rater variability. More importantly, such an automated system can remove a time-consuming, tedious and repetitive task from clinicians and thereby helps to reduce workload of clinicians and/or allowing more time for more patient-centered care. Overall, automation can support clinicians to make healthcare decisions more accurate and timely.

Over the past years, deep learning has made significant headway in achieving above human-level performance on well-specified and isolated tasks in computer vision; for instance in the detection of pulmonary tuberculosis from chest X-rays (Lakhani and Sundaram 2017) and diabetic retinopathy (Dai et al. 2021). The success story of deep learning in computer vision is based on the development of novel deep model architectures that learn complex non-linear features alongside exponential improvements in computer hardware. CNNs are the most commonly applied models in computer vision as they preserve local image relations, while absorbing the feature engineering step into a learning step. This enables capturing critical feature relationships in an image, such as how pixels form a geometric shape (Chai et al. 2021). Thus, important hierarchical feature representations within the data can be discovered without laborious hand-crafting of features (Ker et al. 2018). Consequently, deep learning is the ideal approach for a well circumscribed learning problem such as ROCF scoring and other similar visual-based neuropsychological tests (see examples below). Further, for such tests the ground truth is less controversial than for other problems in neuropsychology or medicine, such as diagnosing patients or predicting treatment response.

However, unleashing the power of deep learning depends on the availability of large training sets with accurate training labels. We obtained over 20k hand-drawn ROCF images and invested immense time and labor to manually scan the ROCF drawings. This effort was only possible by orchestrating across multiple clinical sites and having funds available to hire the work force to accomplish the digitization. Our approach highlights the need to intensify national and international effort of the digitalization of health record data. Although electronic health

records (EHR) are increasingly available (e.g. quadrupled in the US from 2007 to 2012 (Hsiao, Hing, and Ashman 2014)), challenges remain in terms of handling, processing and moving such big data. Improved data accessibility and better consensus in organization and sharing could simplify and foster similar approaches in neuropsychology and medicine.

Another important difficulty is that data typically comes from various sources and thus exhibits large heterogeneity in terms of quality, size, format of the images, and crucially the quality of the labeled dataset. High-quality labeled and annotated datasets are the foundation of a successful computer vision system. A novelty of our approach is that we have trained a large pool of human internet workers (crowdsourced human intelligence) to score ROCFs drawings guided by our self-developed interactive web application. Each element of the figure was scored by several human workers (13 workers on average per figure). To derive ground truth scores, we took the median score, as it has the advantage of being robust to outliers.To further ensure high-quality data annotation, we identified and excluded crowdsourcing participants that have a high level of disagreement (>20% disagreement) with this rating from trained clinicians, who carefully scored manually a subset of the data in the same interactive web application.

Importantly our two-fold innovative approach that combines the digitization of neuropsychological tests and the high-quality scoring using crowdsourced intelligence can provide a roadmap for how AI can be leveraged in neuropsychological testing more generally as it can be easily adapted and applied to various other neuropsychological tests (e.g. Clock Drawing Test (Freedman et al. 1994), Taylor Complex Figure Test (Awad et al. 2004), Hamasch 5-point test (Regard, Strauss, and Knapp 1982)).

Another important prerequisite of using AI for the automated scoring of neuropsychological tests is the availability of training data that is sufficiently diverse and obtains sufficient examples from all possible scores. Recent examples in computer vision have demonstrated that insufficient diversity in training data can lead to biases and adverse consequences (Buolamwini and Gebru 23--24 Feb 2018; Ensign et al. 2017). Given that our training data set included data from children and a representative sample of patients (see Figure 1B), exhibiting a range of neurological disorders, drawings were frequently extremely dissimilar to the target figure. Importantly, our scoring system delivers accurate scores even in cases where drawings are distorted (e.g. caused by fine motor impairments) or omissions (e.g. due to visuospatial deficits), which are typical impairment patterns in neurological patients. In addition, our scoring system is

robust against different semantic transformations (e.g. rotations, perspective change, brightness) which are likely to occur in real life scenarios, when the examiner takes a photo of the ROCF drawing, which naturally will lead to varying viewpoints and illumination conditions. Thus, this robustness is a pivotal prerequisite for any potential clinical utility.

To improve usability of our system and guarantee clinical utility, we have integrated our model into a tablet- (and smartphone-) based application, in which users can take a photo (or upload an image) of an ROCF drawing and the application instantly provides a total score. Importantly, the automated scoring system also maintains explainability by providing visualization of detailed score breakdowns (i.e. item-specific scores), which is a highly valued property of AI-assisted medical decision making and key to help interpret the score and communicate them to the patient (see supplementary Figure S5). The scoring is completely reproducible and thus facilitates valid large-scale comparisons of ROCF data, which enable population-based cognitive screening and (assisted) self-assessments.

In summary, we have created a highly robust and automated AI-based scoring tool, which provides unbiased, reproducible and immediate scores for the ROCF test in research and clinical environments. Our current machine-based automated quantitative scoring system outperforms both amateur raters and clinicians.
The present findings demonstrate that automated scoring systems can be developed to advance the quality of neuropsychological assessments by diminishing reliance on subjective ratings and efficiently improving scoring in terms of time and costs. Our innovative approach can be translated to many different neuropsychological tests, thus providing a roadmap for paving the way to AI-guided neuropsychology.

Materials and Methods

Data

For our experiments, we used a data set of 20225 hand-drawn ROCF images collected from research and clinical settings in various countries spanning the complete range of scores (Figure 1F). Our convergence analysis (see supplementary analysis), suggests that the error converges when approaching 20’000 digitized ROCFs. The data is collected from various populations regarding age and diagnostic status (healthy or with neurological and /or psychiatric

disorder), shown respectively in Figure 1E and Figure 1D. The demographics of the participants and some example images are shown respectively in Figure 1B and Figure 1C. For a subset of the figures (4030), the clinician's scores were available. For each figure only one clinician rating is available. The clinicians ratings derive from six different international clinics (University Hospital Zurich, Switzerland; University Children's Hospital Zurich, Switzerland; BioCruces Health Research Institute, Spain; I.R.C.C.S. Istituto Auxologico Italiano, Ospedale San Giuseppe, Italy; Huashan Hospital, China; University Hospital Magdeburg, Germany). All collaborators have informed consent and/or data usage agreements for the recorded drawings from the participants.

The ROCF images in our dataset have varying resolutions, ranging from 100 x 140 pixels to 3500 x 5300 pixels. Since our models are trained on a fixed input resolution, we investigated the effect of different resolutions on the model performance measured in terms of MAE. To that end, we trained the multilabel classification model with and without data augmentation for inputs of size 78x100, 116x150, 232x300 and 348x450. In principle, when using smaller images, then more information is lost due to resizing, while, on the other hand, a resolution which is too large requires bigger models due to increased complexity in the underlying distribution. In addition, using our data augmentation pipeline with small images might negatively affect the performance since the interpolation techniques used in the semantic transformations like rotations, potentially leads to an additional loss of information. We observed this in our experiments which showed that inputs of size 232x300 yielded the best performance, both for the model with and without data augmentation (Figure 2E). Thus, all subsequent analyses were performed with images of size 232x300.

Convergence Analysis/ Effect of training set size

To get an estimate on the number of ROCF needed to train our models we conducted a convergence analysis. That is, for a fixed test set with 4045 samples, we used different fractions of the remaining dataset to train the multilabel classification model, ending up with training set sizes of 1000, 2000, 3000, 4000, 6000, 8000, 12000, and 16000 samples. By evaluating the performance (i.e. MAE) of the resulting models on the fixed test set, we determined what amount of data is required for the model performance to converge.

Crowdsourced human intelligence for ground truth score

To reach high accuracy in predicting individual sample scores of the ROCFs, it is imperative that the scores of the training set are based on a systematic scheme with as little human bias as possible influencing the score. However, our analysis (see results section) and previous work (Canham, Smith, and Tyrrell 2000b) suggested that the scoring conducted by clinicians may not be consistent, because the clinicians may be unwittingly influenced by the interaction with the patient/participant or by the clinicians factor (e.g. motivation and fatigue). For this reason, we have harnessed a large pool (~5000) of human workers (crowdsourced human intelligence) who scored ROCFs, guided by our self-developed interactive web application (see supplementary Fig. S2). In this web application, the worker was first trained by guided instructions and examples. Subsequently, the worker rated each area of real ROCF drawings separately to guarantee a systematic approach.

For each of the 18 elements in the figure, participants were asked to answer three binary questions: Is the item visible & recognizable? Is the item in the right place? Is the item drawn correctly? This corresponds to the original Osterrieth scoring system (J. E. Meyers, Bayless, and Meyers 1996) (see supplementary Figure S3). The final assessment consisted of answers to these questions for each of the 18 items and enabled calculation of item-wise scores (possible scores: 0, 0.5, 1, 2), which enabled to compute the total score for each image (i.e. sum over all 18 item-wise scores: range total score: 0-36).

Since participants are paid according to how many images they score, there is a danger of collecting unreliable scores when participants rush through too quickly. In order to avoid taking such assessments into account, 600 images have also been carefully scored manually by trained clinicians at the neuropsychological unit of the University Hospital Zurich (in the same interactive web application). We ensured that each crowdsourcing worker rated at least two of these 600 images, which results in 108 ratings (2 figures * 18 items * 3 questions) to compute a level of disagreement. The assessments of crowdsourcing participants that have a high level of disagreement (>20% disagreement) with this rating from clinicians are considered cheaters and are excluded from the data set. After this data cleansing, there remained an average of 13.08 crowdsource-based assessments/scores per figure. In order to use this information for training an algorithm, we require one ground truth score for each image. We assume that the scores approximately follow a normal distribution centered around the true score. With this in mind, we have obtained the ground truth for each drawing by computing the median for each item in the

figure, and then summed up the medians to get the total score for the drawing in question. Taking the median has the advantage of being robust to outliers.

Human MSE

As described in the previous section there are multiple independent scores (from different crowdsourcing workers) available for each image, which often differ substantially. It frequently happens that two people scoring the same image produce different scores. In order to quantify the disagreement among the crowdsourcing participants, we calculated the empirical standard deviation of the scores for each image. With *s1,...,sk* referring to the *k* crowdsourcing scores for a given image and $\bar{s}$ to their mean, the empirical standard deviation is calculated as

$$SD_{empirical} = \sqrt{\frac{1}{k-1}\sum_{i=1}^{k}(s_i - \bar{s})^2}$$

The mean empirical standard deviation is 3.25. Using as ground truth the sum of item score medians, we can also calculate a *human MSE*. This metric describes how close one assessment is to this ground truth on average and lets us make a statement on the difficulty of the task. Reusing above notation, we denote the k crowdsourcing scores for a given image by *s1,…,sk*. Let $\tilde{s}$ be their median. For an image or item, we define the *human MSE* as

$$MSE_{human} = \frac{1}{k}\sum_{i=1}^{k}(s_i - \tilde{s}_i)^2.$$

Similarly, we define the human MAE as

$$MAE_{human} = \frac{1}{k}\sum_{i=1}^{k}|s_i - \tilde{s}_i|.$$

Based on clinician ratings, we define both $MSE_{clinician}$ and $MAE_{clinician}$ in similar fashion. These metrics are used to compare individual crowdsourced scores, clinician scores and model-based scores.

Convolutional Neural Network (CNN)

The overall architecture is shown in Figure 2A. The architecture consists of a shared feature extractor network and eighteen parallel decision heads, one for each item score. The shared feature extractor is composed of three blocks, introducing several model hyperparameters: Each block consists of a 3x3 convolutional layer, batch normalization, 3x3 convolutional layer, batch normalization and max pooling with 2x2 stride. The ReLU activation function is applied after each batch normalization and dropout is applied after max pooling. The output of each block is a representation with 64, 128 and 256 channels, respectively. Each of the eighteen per-item networks is applied on this representation. These consist of a similar feature extractor block, but with global max pooling instead, which outputs a 512-dimensional representation that is subsequently fed into a fully connected layer, that outputs a 1024-dimensional vector, followed by batch normalization, ReLU, dropout and an output fully connected layer.

Two variations of this network were implemented: a regression-based network and a multilabel classification network. They differ in the output layer and in the loss function used during training. The regression-based network outputs a single number, while the multilabel classification network outputs class probabilities, each of which corresponds to one of the four different item scores {0, 0.5, 1.0, 2.0}. Secondly, the choice of loss functions also incurs different dynamics during optimization. The MSE loss used for the regression model penalizes smaller errors less than big errors, e.g., for an item with score 0.5, predicting the score 1.0 is penalized less than when the model predicts 2.0. In other words the MSE loss naturally respects the ordering in the output space. This is in contrast to the multilabel classification model for which a cross entropy loss was used for each item, which does not differentiate between different classes in terms of “how wrong” the model is: in the example above, 1.0 is considered equally wrong as 2.0.

Data augmentation

In a further study, we investigated the effect of data augmentation during training from two perspectives. Firstly, data augmentation is a standard technique to prevent machine learning models from overfitting to training data. The intuition is that enriching the dataset with perturbed versions of the original images leads to better generalization by enriching the dataset with additional and more diverse training set. Secondly, data augmentation can also help in making models more robust against semantic transformations like rotations or changes in perspective.

These transformations are particularly relevant for the present application since users in real-life are likely to take pictures of drawings which might be slightly rotated or with a slightly tilted perspective. With these intuitions in mind, we randomly transformed drawings during training. Each transformation was a combination of Gaussian blur, a random perspective change and a rotation with angles chosen randomly between -10° and 10°.

Training & Validation

To train and evaluate our model, we set aside 4045 of the 20225 ROCF drawings as a testing set, corresponding to 20% of the dataset and used the remaining 16180 images for training and validation. Specifically, we used 12944 (80%) drawings for training and 3236 (20%) as a validation set for early stopping.

The networks and training routines were implemented in PyTorch 1.8.1 (Paszke et al. 2019). The training procedure introduces additional hyperparameters: Our main model was trained on drawings of size 232x300, using batches of 16 images. Model parameters were optimized by the Adam optimizer (Kingma and Ba 2014) with the default parameters $\beta 1 = 0.9$, $\beta 2 = 0.999$ and $\varepsilon = 1e\text{-}8$. The initial learning rate was set to 0.01 and decayed every epoch exponentially with decay rate $\gamma = 0.95$. To prevent overfitting dropout rates were set to 0.3 and 0.5 in the convolutional and fully-connected layers, respectively. Additionally, we used the validation split (20%) of the training data for early stopping. Since our dataset is imbalanced and contains a disproportionately high number of high score drawings, we sampled the drawings in a way that results in evenly distributed scores in a single batch. We trained the model for 75 epochs and picked the weights corresponding to the epoch with the smallest validation loss. The multilabel classification model was trained with cross entropy loss applied to each item classifier independently so that the total loss is obtained by summing the individual loss terms. The regression model was trained in an analogous manner except that we used the MSE for each item score, instead of the cross entropy loss.

Performance evaluation

We evaluated the overall performance of the final model on the test set based on MSE, MAE and R-squared. Even though item-specific scores are discrete and we also implement a classification-based approach, we chose not to rely on accuracy to assess performance. Because Accuracy penalizes errors equally, minor differences in scores can lead to arbitrarily

small Accuracy, or even worse scores than a model that leads to bigger differences on average. MAE gives a better picture of how accurate the model is.

Performance metrics (MSE, MAE, R-squared) were examined for the total score and additionally for each of the 18 items individually. To probe the performance of the model across a wide range of scores, we also obtained MSE and MAE for each total score. Given predictions $\hat{y}$, true scores y and the average of true scores $\bar{y}$, MSE, MAE and R-squared are defined, respectively, as

$$MSE = \frac{1}{N} \sum_{i}^{N} (y_i - \hat{y}_i)^2,$$

$$MAE = \frac{1}{N} \sum_{i}^{N} |y_i - \hat{y}_i|, \text{ and}$$

$$R^2 = 1 - \sum_{i}^{N} (y_i - \hat{y}_i)^2 / \sum_{i}^{N} (y_i - \bar{y})^2.$$

Test-time augmentation

Test-time augmentation (TTA) was performed to promote robustness at test-time to slight variations in input data. First, we experimented with the same augmentations that we used in our data augmentation training policy, namely randomly applying perspective change, resizing the image and applying rotations. In addition to the original image, we fed four to seven randomly transformed drawings (see description of data augmentation for the possible transformations) to the model and averaged the output logits to obtain a prediction. In addition, we experimented with approaches exclusively applying either random rotations or random perspective change. None of these random procedures improved the model performance. Nonetheless, performance improvement was achieved by applying deterministic rotations along positive and negative angles on the unit circle. We performed five rotations, uniformly distributed between -2° and 2°. For this choice, computational restrictions for the model's intended use-case were considered. Specifically, time complexity increases linearly with the number of augmentations. In keeping a small set of augmentations, we believe that small perturbations to rotation present an obvious correspondence to the variations the model will be exposed to in a real application.

Final model selection

Both model and training procedure contain various hyperparameters and even optional building blocks such as using data augmentation (DA) during training, as well as applying test-time-augmentation (TTA) during inference. After optimizing the hyperparameters of each model independently, all of the possible variants that emerge from applying DA and TTA were explored on both classification and regression models. Therefore, the space of possible model variants is spanned by the non-augmented model, the model trained with data augmentation, the model applying test-time-augmentation during inference, as well as the model variant applying both aforementioned building blocks.

To take advantage of particularities in both models, a mixture of the best performing regression and classification models was obtained. In this model, the per-item performances of both models are compared on the held-out validation set. During inference, for each item, we output the prediction of the best performing model. Therefore, on a figure-wide scale, the prediction of the combined model uses both classification and regression models concurrently, while not combining the models' per-item predictions.

Robustness Analysis

We investigated the robustness of our model against different semantic transformations which are likely to occur in real life scenarios. These were rotations, perspective changes, and changes to brightness and contrast. To assess the robustness against these transformations, we transformed images from our test set with different transformation parameters. For rotations, we rotated drawings with angles chosen randomly from increasingly higher orders, both clockwise and counterclockwise. That is, we sampled angles between 0° and 5°, between 5° and 10° and so on, up to 40° to 45°. The second transform we investigated were changes in perspective as these are also likely to occur, for example when photos of drawings are taken with a camera in a tilted position. The degree of perspective change was guided by a parameter between 0 and 1, where 0 corresponds to the original image and 1 corresponds to an extreme change in perspective, making the image almost completely unrecognizable. Furthermore, due to changes in light conditions, an image might appear brighter or with a different contrast. We used brightness parameters between 0.1 and 2.0, where values below 1.0 correspond to a darkening of the image and values above 1.0 correspond to a brighter image. Similarly, for contrast, we varied the contrast parameter between 0.1 and 2.0, with values above (below) 1.0 corresponding to high (low) contrast images.

Supplementary Material

Convergence Analysis

To estimate the number of ROCF needed to train our models a convergence analysis was conducted. With ~3000 images, we obtain diminishing mean MAEs. After ~8000 images, no substantial improvements could be achieved by including more data, as can be seen from the progression plot in Figure 2D. From this, we conclude that the acquired dataset is rich and diverse enough to obtain a powerful machine learning model for the task at hand. In addition, this opens up an avenue for future research in that improvements to the model performance are likely to be achieved via algorithmic improvements, rather than via data-centric efforts.

	MAE	**MSE**	**R^2**
Classifier TTA	1.259 [1.214, 1.308]	3.902 [3.598, 4.268]	0.963
Classifier NA	1.228 [1.185, 1.275]	3.700 [3.418, 4.042]	0.965
Regressor DA	1.162 [1.122, 1.204]	3.174 [2.958, 3.425]	0.970
Classifier DA	1.161 [1.118, 1.206]	3.368 [3.105, 3.743]	0.968
Regressor DA + TTA	1.152 [1.110, 1.194]	3.152 [2.928, 3.399]	0.970
Classifier DA + TTA	1.149 [1.106, 1.194]	3.315 [3.061, 3.646]	0.969
Regressor TTA	1.139 [1.099, 1.181]	3.074 [2.850, 3.335]	0.971
Regressor NA	1.138 [1.098, 1.180]	3.078 [2.850, 3,354]	0.971
Final Model (Best Regressor + Best Classifier)	1.113 [1.074, 1.156]	3.000 [2.775, 3,262]	0.972

Table S1. The MAE for all model variants. NA: non-augmented, DA: data augmentation is performed during training, TTA: test-time augmentation. The 95% confidence interval is shown in square brackets.

	MAE	MSE	R^2		MAE	MSE	R^2
Item 1	0.104	0.095	0.800	**Item 10**	0.051	0.048	0.950
Item 2	0.170	0.166	0.587	**Item 11**	0.081	0.072	0.853
Item 3	0.148	0.151	0.778	**Item 12**	0.101	0.092	0.869
Item 4	0.127	0.132	0.730	**Item 13**	0.116	0.107	0.778
Item 5	0.125	0.132	0.744	**Item 14**	0.094	0.092	0.839
Item 6	0.148	0.131	0.802	**Item 15**	0.073	0.090	0.899
Item 7	0.063	0.073	0.923	**Item 16**	0.094	0.101	0.862
Item 8	0.105	0.090	0.860	**Item 17**	0.091	0.078	0.869
Item 9	0.091	0.084	0.893	**Item 18**	0.114	0.108	0.814
				Total score	1.113	3.000	0.972

Table S2. Per-item and total performance estimates for the final model. MAE, MSE and R^2 are estimated directly from the estimated scores. Multiclass Accuracy is obtained after the discretization of predicted scores. Accuracy is not evaluated for the total score, since it consists of the sum of item scores.

Score Interval	MSE	MAE	Score Interval	MSE	MAE	Score Interval	MSE	MAE	Score Interval	MSE	MAE
(0 - 0.5) (n=22)	3.72	1.52	(10 - 10.5) (n=104)	5.52	1.79	(19 - 19.5) (n=42)	3.07	1.25	(28 - 28.5) (n=130)	3.28	1.27
(1 - 1.5) (n=15)	2.78	1.17	(11 - 11.5) (n=33)	3.17	1.41	(20 - 20.5) (n=131)	4.90	1.56	(29 - 29.5) (n=92)	3.11	1.28
(2 - 2.5) (n=52)	3.15	1.41	(12 - 12.5) (n=103)	5.77	1.74	(21 - 21.5) (n=45)	4.41	1.30	(30 - 30.5) (n=155)	3.26	1.28
(3 - 3.5) (n=26)	3.93	1.52	(13 - 13.5) (n=28)	2.77	1.36	(22 - 22.5) (n=105)	3.75	1.44	(31 - 31.5) (n=134)	3.33	1.38
(4 - 4.5) (n=62)	3.66	1.52	(14 - 14.5) (n=108)	4.82	1.56	(23 - 23.5) (n=75)	3.26	1.05	(32 - 32.5) (n=167)	2.67	1.15
(5 - 5.5) (n=32)	7.23	2.09	(15 - 15.5) (n=48)	3.70	1.52	(24 - 24.5) (n=120)	3.43	1.25	(33 - 33.5) (n=196)	1.73	0.96
(6 - 6.5) (n=92)	5.27	1.61	(16 - 16.5) (n=141)	5.17	1.51	(25 - 25.5) (n=75)	2.45	1.03	(34 - 34.5) (n=299)	1.21	0.77
(7 - 7.5) (n=21)	6.96	1.74	(17 - 17.5) (n=54)	4.40	1.56	(26 - 26.5) (n=149)	4.32	1.44	(35 - 35.5) (n=295)	0.88	0.58
(8 - 8.5) (n=88)	4.79	1.52	(18 - 18.5) (n=115)	6.03	1.51	(27 - 27.5) (n=93)	2.82	1.11	(36) (n=564)	0.45	0.20
(9 - 9.5) (n=34)	2.21	1.13									

Table S3. Performance per total score interval. Thirty seven intervals were evaluated, across the whole range of scores. We evaluated MAE and MSE for the total score within each interval.

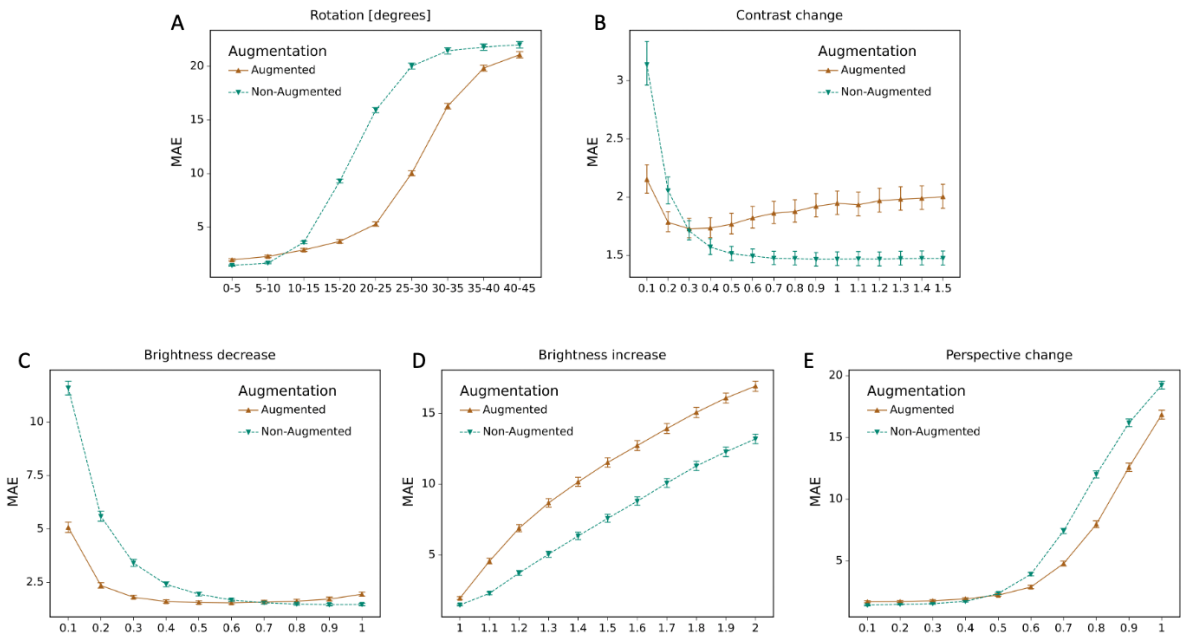

Supplementary Figure S1. Effect of data augmentation

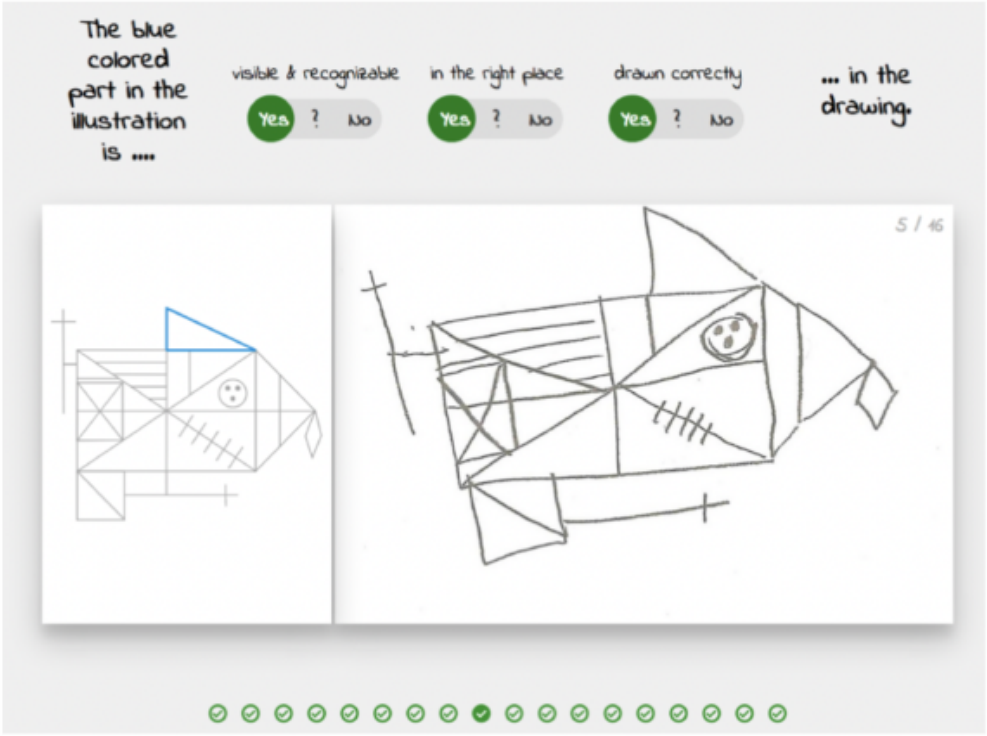

Supplementary Figure S2. The graphical user interface of the crowdsourcing application

Unit Correct	Placed Properly	2
Unit Distorted, Incomplete but Recognisable	Placed Poorly	1
	Placed Properly	1
	Placed Poorly	1/2
Absent or Unrecognisable		0

Supplementary Figure S3: Original scoring system according to Osterrieth

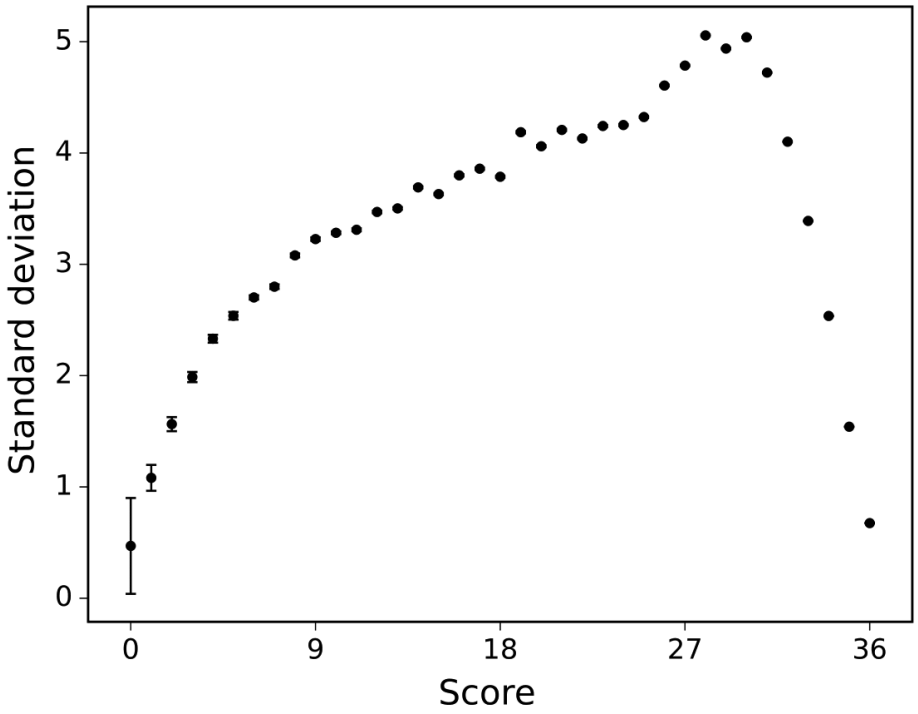

Supplementary Figure S4: The standard deviation of the human raters is displayed across differently scored drawings.

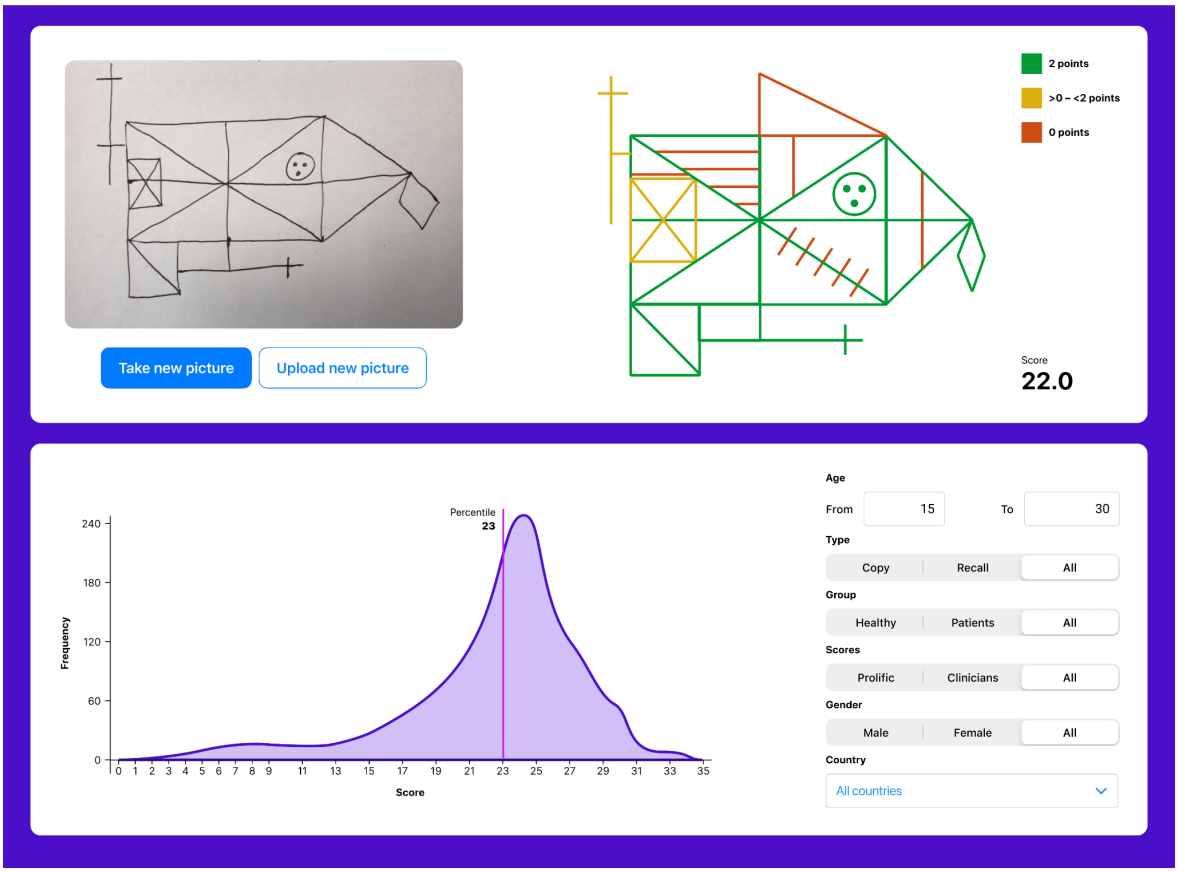

Supplementary Figure S5: The user interface for the tablet- (and smartphone-) based application. The application enables explainability by providing a score for each individual item. Furthermore, the total score is displayed. The user can also compare the individual with a choosable norm population.

References

Alladi, Suvarna, Robert Arnold, Joanna Mitchell, Peter J. Nestor, and John R. Hodges. 2006. "Mild Cognitive Impairment: Applicability of Research Criteria in a Memory Clinic and Characterization of Cognitive Profile." *Psychological Medicine*. https://doi.org/10.1017/s0033291705006744.

Amiri, Mina, Rupert Brooks, Bahareh Behboodi, and Hassan Rivaz. 2020. "Two-Stage Ultrasound Image Segmentation Using U-Net and Test Time Augmentation." *International Journal of Computer Assisted Radiology and Surgery* 15 (6): 981–88.

Awad, Nesrine, Maria Tsiakas, Michèle Gagnon, Valérie B. Mertens, Eva Hill, and Claude Messier. 2004. "Explicit and Objective Scoring Criteria for the Taylor Complex Figure Test." *Journal of Clinical and Experimental Neuropsychology* 26 (3): 405–15.

Berry, David T. R., Rebecca S. Allen, and Frederick A. Schmitt. 1991. "Rey-Osterrieth Complex Figure: Psychometric Characteristics in a Geriatric Sample." *Clinical Neuropsychologist*. https://doi.org/10.1080/13854049108403298.

Buolamwini, Joy, and Timnit Gebru. 23--24 Feb 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, edited by Sorelle A. Friedler and Christo Wilson, 81:77–91. Proceedings of Machine Learning Research. PMLR.

Canham, R. O., S. L. Smith, and A. M. Tyrrell. 2000a. "Automated Scoring of a Neuropsychological Test: The Rey Osterrieth Complex Figure." *Proceedings of the 26th Euromicro Conference. EUROMICRO 2000. Informatics: Inventing the Future*. https://doi.org/10.1109/eurmic.2000.874519.

———. 2000b. "Automated Scoring of a Neuropsychological Test: The Rey Osterrieth Complex Figure." *Proceedings of the 26th Euromicro Conference. EUROMICRO 2000. Informatics: Inventing the Future*. https://doi.org/10.1109/eurmic.2000.874519.

Chai, Junyi, Hao Zeng, Anming Li, and Eric W. T. Ngai. 2021. "Deep Learning in Computer Vision: A Critical Review of Emerging Techniques and Application Scenarios." *Machine Learning with Applications* 6 (December): 100134.

Dai, Ling, Liang Wu, Huating Li, Chun Cai, Qiang Wu, Hongyu Kong, Ruhan Liu, et al. 2021. "A Deep Learning System for Detecting Diabetic Retinopathy across the Disease Spectrum." *Nature Communications* 12 (1): 3242.

Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2017. "Runaway Feedback Loops in Predictive Policing." *arXiv [cs.CY]*. arXiv. http://arxiv.org/abs/1706.09847.

Fastenau, Philip S., Jeanne M. Bennett, and Natalie L. Denburg. 1996. "Application of Psychometric Standards to Scoring System Evaluation: Is 'New' Necessarily 'Improved'?" *Journal of Clinical and Experimental Neuropsychology*. https://doi.org/10.1080/01688639608409003.

Franzen, Michael D. 2000. "Validity as Applied to Neuropsychological Assessment." *Reliability and Validity in Neuropsychological Assessment*. https://doi.org/10.1007/978-1-4757-3224-5_5.

Freedman, Morris, Larry Leach, Edith Kaplan, Gordon Winocur, Kenneth Shulman, Dean C. Delis, and Others. 1994. *Clock Drawing: A Neuropsychological Analysis*. Oxford University Press, USA.

Groth-Marnat, Gary. 2000. *Neuropsychological Assessment in Clinical Practice: A Guide to Test Interpretation and Integration*. Wiley.

Harbi, Zainab, Yulia Hicks, and Rossitza Setchi. 2016. "Clock Drawing Test Digit Recognition Using Static and Dynamic Features." *Procedia Computer Science*.

https://doi.org/10.1016/j.procs.2016.08.166.
Hsiao, Chun-Ju, Esther Hing, and Jill Ashman. 2014. “Trends in Electronic Health Record System Use among Office-Based Physicians: United States, 2007-2012.” *National Health Statistics Reports*, no. 75 (May): 1–18.
Huygelier, Hanne, Margaret Jane Moore, Nele Demeyere, and Céline R. Gillebert. 2020. “Non-Spatial Impairments Affect False-Positive Neglect Diagnosis Based on Cancellation Tasks.” *Journal of the International Neuropsychological Society: JINS* 26 (7): 668–78.
Ker, Justin, Lipo Wang, Jai Rao, and Tchoyoson Lim. 2018. “Deep Learning Applications in Medical Image Analysis.” *IEEE Access: Practical Innovations, Open Solutions* 6: 9375–89.
Kim, Hyungsin, Young Suk Cho, and Ellen Yi-Luen Do. 2010. “Computational Clock Drawing Analysis for Cognitive Impairment Screening.” *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction*. https://doi.org/10.1145/1935701.1935768.
Kingma, Diederik P., and Jimmy Ba. 2014. “Adam: A Method for Stochastic Optimization,” December. https://doi.org/10.48550/arXiv.1412.6980.
Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2017. “ImageNet Classification with Deep Convolutional Neural Networks.” *Communications of the ACM*. https://doi.org/10.1145/3065386.
Lakhani, Paras, and Baskaran Sundaram. 2017. “Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks.” *Radiology* 284 (2): 574–82.
Levine, Andrew J., Eric N. Miller, James T. Becker, Ola A. Selnes, and Bruce A. Cohen. 2004. “Normative Data for Determining Significance of Test–Retest Differences on Eight Common Neuropsychological Instruments.” *The Clinical Neuropsychologist*. https://doi.org/10.1080/1385404049052420.
Meyers, J. E., J. D. Bayless, and K. R. Meyers. 1996. “Rey Complex Figure: Memory Error Patterns and Functional Abilities.” *Applied Neuropsychology* 3 (2): 89–92.
Meyers, John E., and Marie Volbrecht. 1999. “Detection of Malingerers Using the Rey Complex Figure and Recognition Trial.” *Applied Neuropsychology*. https://doi.org/10.1207/s15324826an0604_2.
Moshkov, Nikita, Botond Mathe, Attila Kertesz-Farkas, Reka Hollandi, and Peter Horvath. 2020. “Test-Time Augmentation for Deep Learning-Based Cell Segmentation on Microscopy Images.” *Scientific Reports* 10 (1): 5068.
Nazar, Haris Bin, Momina Moetesum, Shoaib Ehsan, Imran Siddiqi, Khurram Khurshid, Nicole Vincent, and Klaus D. McDonald-Maier. 2017. “Classification of Graphomotor Impressions Using Convolutional Neural Networks: An Application to Automated Neuro-Psychological Screening Tests.” *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*. https://doi.org/10.1109/icdar.2017.78.
Osterrieth, Paul Alexandre. 1944. *Le test de copie d’une figure complexe: Contribution à l'étude de la perception et de la mémoire*. Delachaux et Niestlé.
Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, et al. 2019. “PyTorch: An Imperative Style, High-Performance Deep Learning Library.” In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. https://doi.org/10.48550/arXiv.1912.01703.
Petilli, Marco A., Roberta Daini, Francesca Lea Saibene, and Marco Rabuffetti. 2021. “Automated Scoring for a Tablet-Based Rey Figure Copy Task Differentiates Constructional, Organisational, and Motor Abilities.” *Scientific Reports* 11 (1): 14895.
Rabin, Laura A., William B. Barr, and Leslie A. Burton. 2005. “Assessment Practices of Clinical Neuropsychologists in the United States and Canada: A Survey of INS, NAN, and APA Division 40 Members.” *Archives of Clinical Neuropsychology: The Official Journal of the National Academy of Neuropsychologists* 20 (1): 33–65.

Rabin, Laura A., Emily Paolillo, and William B. Barr. 2016. “Stability in Test-Usage Practices of Clinical Neuropsychologists in the United States and Canada Over a 10-Year Period: A Follow-Up Survey of INS and NAN Members.” *Archives of Clinical Neuropsychology: The Official Journal of the National Academy of Neuropsychologists* 31 (3): 206–30.

Regard, Marianne, Esther Strauss, and Paul Knapp. 1982. “Children’s Production on Verbal and Non-Verbal Fluency Tasks.” *Perceptual and Motor Skills*. https://doi.org/10.2466/pms.1982.55.3.839.

Shin, Min-Sup, Sun-Young Park, Se-Ran Park, Soon-Ho Seol, and Jun Soo Kwon. 2006a. “Clinical and Empirical Applications of the Rey–Osterrieth Complex Figure Test.” *Nature Protocols*. https://doi.org/10.1038/nprot.2006.115.

———. 2006b. “Clinical and Empirical Applications of the Rey–Osterrieth Complex Figure Test.” *Nature Protocols*. https://doi.org/10.1038/nprot.2006.115.

———. 2006c. “Clinical and Empirical Applications of the Rey–Osterrieth Complex Figure Test.” *Nature Protocols* 1 (2): 892–99.

Somerville, J., G. Tremont, and R. A. Stern. 2000. “The Boston Qualitative Scoring System as a Measure of Executive Functioning in Rey-Osterrieth Complex Figure Performance.” *Journal of Clinical and Experimental Neuropsychology* 22 (5): 613–21.

Trautmann, Sebastian, Jürgen Rehm, and Hans‐ulrich Wittchen. 2016. “The Economic Costs of Mental Disorders.” *EMBO Reports*. https://doi.org/10.15252/embr.201642951.

Trojano, Luigi, and Guido Gainotti. 2016. “Drawing Disorders in Alzheimer’s Disease and Other Forms of Dementia.” *Journal of Alzheimer’s Disease*. https://doi.org/10.3233/jad-160009.

Tupler, L. A., K. A. Welsh, Y. Asare-Aboagye, and D. V. Dawson. 1995. “Reliability of the Rey-Osterrieth Complex Figure in Use with Memory-Impaired Patients.” *Journal of Clinical and Experimental Neuropsychology* 17 (4): 566–79.

Vogt, J., H. Kloosterman, S. Vermeent, G. Van Elswijk, R. Dotsch, and B. Schmand. 2019. “Automated Scoring of the Rey-Osterrieth Complex Figure Test Using a Deep-Learning Algorithm.” *Archives of Clinical Neuropsychology*. https://doi.org/10.1093/arclin/acz035.04.

Watkins, Marley W. 2017. “The Reliability of Multidimensional Neuropsychological Measures: From Alpha to Omega.” *The Clinical Neuropsychologist* 31 (6-7): 1113–26.

Webb, Sam S., Margaret Jane Moore, Anna Yamshchikova, Valeska Kozik, Mihaela D. Duta, Irina Voiculescu, and Nele Demeyere. 2021. “Validation of an Automated Scoring Program for a Digital Complex Figure Copy Task within Healthy Aging and Stroke.” *Neuropsychology* 35 (8): 847–62.