



Harnessing AI and mathematical modeling for oral cancer survival prediction across stages

Dharmaraja Selvamuthu¹ · Raghav Aggarwal¹

Received: 9 January 2025 / Accepted: 16 February 2026

© The Author(s), under exclusive licence to Operational Research Society of India 2026

Abstract

Mathematical models are finding increased use in biology and particularly in the field of cancer research. With cancer posing a significant global health challenge due to its high mortality rate and complex progression, predictive models have become essential tools in understanding disease dynamics and personalizing interventions. This study delves into the intersection of mathematical modeling and artificial intelligence (AI) techniques for survival prognosis and risk stratification. It focuses on 1-year, 2-year, 3-year, 4-year, and 5-year survivability of cancer patients under different cancer stages and treatment modalities. This paper reviews neural network models used in survivability predictions related to cancer onset, progression, and treatment outcomes, with a major focus on oral cancer. This research evaluates the impact of key design elements and training strategies. Its goal is to develop a state-of-the-art solution for complex classification problems in cancer progression prediction and survival. The study emphasizes the foundational role of mathematics in achieving these advancements. Relative to commonly reported Cox proportional hazards / Random Survival Forest baselines on comparable SEER cohorts (typically 60–65% accuracy and 65–70% F1) (Kim in Sci Rep 9 (1), 6994, 2019), our best neural configurations reach 72.64–87.86% accuracy across 5-year to 1-year horizons, implying an absolute misclassification reduction of approximately 10–23 percentage points for horizon-specific prognosis.

Keywords Oral cancer prediction · Artificial neural networks · Survival prediction models · Deep learning in oncology · Survival rate prediction · Cancer staging

✉ Dharmaraja Selvamuthu
dharmar@maths.iitd.ac.in

Raghav Aggarwal
mt1230238@iitd.ac.in

¹ Department of Mathematics, Indian Institute of Technology Delhi, New Delhi 110016, India

1 Introduction

Cancer is a leading global cause of death, with approximately 9.7M deaths recorded in 2022 worldwide. Most cancer-related fatalities occur in developing countries. In India, 0.91M and in the US, 0.6M deaths were reported due to cancer in 2022. Additionally, 1,88,438 deaths were reported due to oral cancer worldwide in 2022, out of which 79,979 deaths were reported in India alone, followed by 20,831 in China and 11,100 in the United States, as reported by Bray et al.(2024) [2]. Cancer is a serious health problem around the world that scientists and authorities follow very closely because of its high mortality rate. In the last decades, the pressure of cancer in public health has increased. Efforts to combat cancer include improving treatments, enhancing low-cost interventions, and bolstering patients immune systems. Due to its prevalence, research in the domain is highly extended in all spheres, especially early detection, prediction of disease progression, and optimization of treatment. AI has a significant role in enhancing survivability and early detection of oral cancer. AI algorithms can detect abnormalities in oral tissues, identifying early signs of cancer that may be missed by human observation. Machine learning models trained on large datasets of biopsy results, radiographs, and patient histories can predict the likelihood of cancer development. AI can assist in personalizing treatment plans by evaluating a patients' condition and providing recommendations based on data trends. Additionally, AI-driven systems can monitor post-treatment recovery, ensuring timely interventions if cancer recurrence is detected. With the ability to process vast amounts of data quickly, AI can improve diagnostic accuracy, reduce the time for detection, and increase the survival rates of oral cancer patients. Most of the studies are on limited dataset for Oral cancer. There are not many studies to our knowledge on Oral cancer survivability for 1-year, 2-year, 3-year, 4-year, and 5-year survivability corresponding to cancer stages in the absence of cancer reoccurrence data with time or limited data using artificial neural networks. In this study attempt has been made to build the artificial neural network model and establish the mathematics behind the same for this problem statement. For this study, patient data on smoking, alcohol and behavioral patterns was not available.

For clarity: “Stage 1–Stage 4” in Sect. 3, 4 follows the SEER 4-level ordinal tumor grade/differentiation encoding used in our features (Grade Mapping = 1–4), which is prognostically informative; AJCC/TNM stage variables were not consistently available across the extracted cohort.

This paper is organized as follows: Section 2 reviews mathematical modeling for cancer survivability prediction, Section 3 presents the model description and SEER data analysis, Section 4 discusses results including activation functions and 1-year to 5-year survivability predictions, and Section 5 concludes with findings and future research directions.

2 Literature review

The literature with regard to mathematical modeling of survivability of patients, initiation, proliferation, and metastasis of Oral cancer has been reviewed. Mathematical models to forecast and simulate the growth rate of cancer cells have been reviewed from both deterministic and stochastic considerations. These are not very popular for Oral cancer research. As computational methods progressed, data mining and other machine learning techniques became critical tools toward bettering the diagnosis of and prognosis in cancer treatment. These techniques, especially CNN, ANN, Naive Bayes, Markov chains, Bayesian Network, Gaussian mixture models, LDA, etc., are used particularly for modeling very complex relationships in medical data, providing insights into the patterns and results of the disease much easier. Most of the studies are on limited datasets for Oral cancer.

Paper by Shabana Tabassum et al. [3] reviewed deterministic and stochastic models for cancer growth. While deterministic models like Gompertz and logistic show high accuracy compared to others, their stochastic extensions are necessary to account for variability in tumor dynamics. Among these, the stochastic Gompertz model emerges as the most effective in fitting clinical data. Deterministic models excel in explaining early-stage cancer growth but struggle with long-term predictions. Stochastic models, by incorporating randomness, better align with experimental data, making the stochastic Gompertz model particularly noteworthy for describing cancer cell proliferation.

Kim et al. [1] evaluated the deep learning-based survival prediction model, DeepSurv, a neural network survival analysis model for oral squamous cell carcinoma (SCC) patients. They compared the model with traditional Cox proportional hazards (CPH) and Random Survival Forest (RSF) models using clinical data from 255 patients. DeepSurv outperformed both models, yielding the highest c-index scores of 0.810 and 0.781 for training and testing datasets, respectively, showing improved predictive accuracy as features were added. The study presents DeepSurv as a method that can direct personalized treatment strategies by fitting survival outcomes with high accuracy, accounting for nonlinear interactions among clinical variables.

TENG et al. [4] Dynamic Bayesian model for the study of Breast cancer survival prediction developed a Bayesian dynamic Cox (BDCox) model to predict the 5-year overall survival in breast cancer patients utilizing SEER Cancer Registry. Four feature selection methods were utilized, and Bayesian Model Averaging (BMA) outperformed all others in both internal and external validation. The final BDCox model, with 12 predictors, had excellent generalization ability and achieved high accuracy (C-statistic: 0.802 internally, 0.739 externally). Some predictors displayed time-varying survival associations, consistent with the findings of previous studies.

Gang et al. [5] present a probabilistic approach utilizing a Hidden Markov Model (HMM) to estimate the likelihood of occult nodal metastases in individuals diagnosed with oral tongue squamous cell carcinomas (OTSCCs). The model utilizes MRI to assess macroscopic metastases and tumor stage to quantify the probability of microscopic lymph node involvement. Data from 108 OTSCC patients were analyzed using Markov chain Monte Carlo sampling. The results showed high probabilities of metastasis in lymph node levels I and II, with level II metastasis predicting a

small probability (7.88%) of metastasis in level III 93.79%). The model suggests that lymph node metastasis progresses faster in early stages and slower in later stages. This HMM-based approach can predict nodal metastasis evolution and help in reducing unnecessary invasive procedures.

Rajaguru and Prabhakar [6] compared the performance of the GMM, MLP, and ELM classifiers on the classification of oral cancer using clinical data from 75 patients. The inputs used for classification were based on TNM staging variables: tumor size, node involvement, and metastasis. Among the classifiers used, GMM achieved an average accuracy of 94.18% on all stages, which surpassed MLP at 89.5%. ELM was also attempted but resulted in relatively inferior performance. The results reveal the potential of GMM as a sound tool for accurate oral cancer stage classification, thus enabling improved diagnosis and therapeutic planning.

Weatherspoon et al. [7] studied trends in incidence of oral cavity (OC) and oropharyngeal (OP) cancer in the United States during 2000–2010 by Poisson regression. The study indicated a rising trend of incidence for OP cancers, mainly of the tonsil and the base of the tongue. However, the incidence of OC cancers was decreasing. The highest incidence was noted among White males. More at risk in older age groups and being male, these cancers are often diagnosed by regional stage; the paper emphasizes the distinctiveness of both OC and OP cancers, calling for specially targeted prevention and education campaigns.

Baik et al. [8] evaluated Random Forests-based algorithms for classifying oral premalignant lesions (OPLs) using image cytometry. They introduced an automated nuclear phenotypic score (aNPS) to predict the risk of progression to SCC. The algorithm achieved 80 % accuracy at the cellular level and 75 % at the tissue level, with predictive values of 74 % (positive) and 76 % (negative). The aNPS correlated with histological severity and progression risk, supporting its potential in clinical pathology. Limitations included partial automation and the need for larger validation datasets. This study highlights the promise of machine learning in cancer diagnostics.

3 Model description

The refined and preprocessed dataset serves as a foundation for constructing predictive models aimed at understanding and improving oral cancer outcomes. By integrating demographic, clinical, and treatment-related features, this study employs advanced machine learning techniques to uncover patterns and insights that may guide personalized treatment strategies and enhance survivability predictions.

3.1 Data mining

This study utilized cancer incidence data sourced from the publicly accessible National Cancer Institute (NCI) Surveillance, Epidemiology, and End Results (SEER) Research Database, covering 12 registries and submitted in November 2023 (spanning 1992–2021). From a total dataset comprising approximately 4.9 million records of various cancer patients, a subset of 95,338 records pertaining to oral cancer patients was identified. This subset included cases with primary cancer sites such

as the floor of the mouth, gum and other mouth regions, hypopharynx, lip, oropharynx, tongue, and tonsil.

A rigorous data preprocessing pipeline was implemented to ensure data quality and relevance. After removing duplicate entries, incomplete records, and cases lacking sufficient information related to the identified cancer sites, a final dataset of 10,978 unique patient records was prepared, focused on diagnosis years 2010 to 2015.

Given the diversity of the dataset, key clinical and demographic features known to influence oral cancer progression were identified, ensuring the data comprehensively represents the patient population and treatment modalities.

Seventeen independent features were selected based on their relevance and availability. These include patient-specific data such as Patient ID, year of diagnosis, sex, age, and ethnicity/race. Clinical variables such as cancer staging, diagnostic confirmation, histology codes, tumor size, lymph node involvement, and treatment modalities including surgery, chemotherapy, and radiotherapy were incorporated to reflect comprehensive patient profiles.

Outcome definition, censoring, and class balance. SEER is a time-to-event registry; for clinically interpretable, horizon-specific prognosis we define fixed-horizon labels for $t \in \{1, 2, 3, 4, 5\}$ years: $y_i^{(t)} = \mathbb{I}\{\text{patient } i \text{ survives at least } t \text{ years from diagnosis}\}$ (0 if death occurs before t , 1 otherwise). Records with missing/unknown vital status or follow-up were removed in preprocessing. With diagnosis years 2010–2015 and follow-up through 2021, administrative censoring for $t \leq 5$ years is substantially reduced; remaining censoring bias is addressed in Future Work through survival-analysis objectives and sensitivity checks. Table 1 reports per-horizon prevalence and the class weights used (inverse prevalence) for stable training.

Why fixed-horizon classification (and how we handle remaining censoring). Fixed-horizon outcomes are widely used in clinical survivability reporting because they align with milestone-based planning (e.g., 1-year near-term surveillance versus 5-year survivorship planning). Our goal is horizon-specific prognosis and risk stratification rather than estimation of an entire survival curve. To avoid labeling ambiguity, if a patient is alive but last known follow-up occurs before horizon t , then survival status at t is not observable; such records are excluded rather than treated as negative. Concretely, letting T_i denote observed time from diagnosis to last contact (or death) and δ_i denote the event indicator (1 if death observed, 0 otherwise), a label $y_i^{(t)}$ is only assigned when either $\delta_i = 1$ and $T_i < t$ (negative) or when $T_i \geq t$ (positive). Patients with $\delta_i = 0$ and $T_i < t$ are right-censored before the horizon and are removed for that horizon. Because our cohort window has follow-up through 2021, most records satisfy $T_i \geq 5$ years, reducing the impact of this exclusion, but any remaining bias is acknowledged below and addressed as future work.

Table 1 Per-horizon outcome prevalence and class-weight summary (N=10,978)

Horizon	Positives	Negatives	Prevalence (%)	w_1	w_0
1-year	9,888	1,090	90.07	0.555	5.036
2-year	9,403	1,575	85.64	0.584	3.485
3-year	8,860	2,118	80.70	0.619	2.591
4-year	8,314	2,664	75.73	0.660	2.061
5-year	6,739	4,239	61.39	0.814	1.295

Here $w_1 = \frac{N}{2N_1}$ and $w_0 = \frac{N}{2N_0}$ are horizon-wise class weights for positive and negative classes, respectively (with N_1 positives and N_0 negatives).

Time ordering, leakage prevention, and treatment-variable scope. All model inputs are baseline variables available at diagnosis or during first-course treatment reporting in SEER (demographics, tumor characteristics, and first-course treatment indicators). We explicitly exclude any survival-time derived variables (e.g., survival months, cause-specific survival time) from the feature set. Treatment covariates are interpreted as first-course indicators recorded around the time of diagnosis and therefore precede the fixed-horizon outcomes. We report a temporal diagnosis-year split (2010–2014 train, 2015 test) to reflect prospective deployment and prevent temporal leakage, while still allowing year-of-diagnosis to capture real-world era effects (changes in practice patterns) that would also be present at deployment time.

The distribution of Oral cancer cases across anatomical sites highlighted tongue cancer as predominant, followed by tonsil and gum cancers. This distribution underscores the need for targeted approaches and provides context for subsequent analysis (Fig. 1).

Racial demographics reveal that 76.4% of cases were reported in non-Hispanic White individuals, indicating possible demographic biases in disease incidence and access to care (Fig. 2).

Age and gender distributions show 70% of patients falling between 50 and 75 years and two-thirds being male; this aligns with epidemiological trends in oral cancer, emphasizing age-related susceptibility and the gender disparity (Fig. 3). Histological data indicates 95.8% squamous cell carcinoma presence, consistent with literature that SCC represents the majority of oral cancers. This reinforces the dataset's clinical relevance.

Cancer spread and lymph node involvement data provide spatial and metastatic insights. Tongue cancers, with the highest spread (42%) and lymph node involve-




Fig. 1 Distribution of oral cancer




Fig. 2 Distribution of races




Fig. 3 Distribution of stages of cancer across different age groups

ment in nearly half the cases, suggest aggressive local disease progression requiring tailored treatment strategies (Fig. 4).

Treatment analysis shows 69.6% surgery uptake and widespread use of radiotherapy (58%) and chemotherapy (59%), with only 17.1% receiving all three modalities. These distributions highlight real-world treatment patterns and their potential impacts on survival (Figs. 5, 6).

3.2 Feature modeling and encoding

Features were grouped into demographic, clinical, and treatment categories. Patient ID was excluded from modeling as a unique identifier. Year of diagnosis was treated continuously to capture temporal trends, while categorical data such as sex and ethnicity were one-hot encoded for unbiased treatment.

Age was discretized into four non-overlapping categories, capturing nonlinear effects of aging on disease progression:

- 0–20: years for rare, early onsets.
- 20–55: years capturing younger adults with distinct risk profiles.
- 55–75: years as the modal age range with peak incidence.
- 75+: years covering older adults with comorbidities impacting treatment.

Cancer stage was treated as an ordinal variable integral to survivability prediction. Table 2 catalogs case distribution across stages and sites, providing necessary context for model development.

Diagnostic confirmation and histology codes were one-hot encoded to model diverse diagnostic approaches and tumor pathologies. (Table 3) Tumor size, retained




Fig. 4 Site of cancer for different cancer types




Fig. 5 Distribution of surgery performed for patients in the dataset




Fig. 6 Treatment modality distribution

as continuous, reflected progression severity while lymph node involvement was binary encoded as present or absent.

Treatment features modeled surgery, chemotherapy, and radiotherapy as binary variables, accounting for combination effects, informed by observed treatment patterns (Figs. 5, 6).

Table 2 Stages of cancer at different sites

Cancer site	Stage-1	Stage-2	Stage-3	Stage-4	Total
Floor of Mouth	129	446	129	6	710
Gum and Other Mouth	405	888	279	20	1592
Hypopharynx	25	256	236	9	526
Lip	317	333	83	2	735
Oropharynx	20	157	128	11	316
Tongue	818	2285	1535	47	4685
Tonsil	108	917	1339	50	2414
Total	1822	5282	3729	145	10978

Table 3 Descriptions and unique values of variables

Variable name	Description	Unique Value Count
Patient ID	Patient's unique identifier	10,978
Year of diagnosis	Year diagnosis was made	6
Age Mapping	Discretized age categories	4
Site record mapping	Primary tumor sites	15
Sex Mapping	Gender categories	2
Grade Mapping	Histologic grading	4
AYA Site mapping	Adolescent/Young Adult site groups	8
Diagnosis Mapping	Diagnosis codes	10
Tumor Size Over Time Recode	Tumor size continuum	Continuous
Histological recode broad grouping	Broad tumor classifications	12
Surgery Mapping	Surgery codes	5
Radiotherapy mapping	Radiotherapy presence	2
Chemotherapy mapping	Chemotherapy presence	2
Race Mapping	Ethnicity categories	7
Lymph mapping	Lymph node presence	2
SEER historic stage A	SEER staging system	3

All categorical variables, except tumor size and year of diagnosis, were one-hot encoded to avoid assuming ordinal relationships. Continuous variables were normalized to ensure homogeneity in model training.

This structured feature engineering ensured comprehensive capture of patient, tumor, and treatment factors essential for predictive modeling. Encoding preserved feature integrity and facilitated effective machine learning for survival analysis.

3.3 Model architecture and data preprocessing

3.3.1 Overview of dataset

The first round of training for the ANN was performed on the SEER dataset, a comprehensive resource of cancer incidence and survival data. This dataset is specifically compiled for predicting multi-year survivability (1-year to 5-year survivability) of

oral cancer patients across different cancer stages, making it ideal for our survival prediction objectives. After removing missing, duplicate, and inconsistent records, 10,978 complete entries remained suitable for training.

3.3.2 Feature engineering and preprocessing

The feature set was specifically selected to align with clinical oral cancer prognostic factors: categorical variables such as age, cancer site, sex, race, diagnosis type, stage, surgery, chemotherapy, radiotherapy, and lymph node involvement, along with numerical variables such as year of diagnosis and tumor size. These features directly correspond to established TNM staging criteria and treatment protocols used in oral cancer management. Categorical features were one-hot encoded to ensure compatibility with the ANN's numerical input requirements. Numerical features were standardized using z-scores to ensure optimal neural network interpretation of feature values.

To reduce temporal leakage, all covariates were restricted to baseline (diagnosis/first-course) descriptors and survival-time variables were never used as inputs.

3.4 Artificial neural network (ANN) architecture

3.4.1 Overview of ANN Design

The ANN was specifically designed for oral cancer survival prediction with multiple densely connected layers with Softplus activation functions for hidden units, enabling smooth, non-linear transformations crucial for capturing complex interactions between clinical variables affecting survival outcomes. Dropout regularization, applied at a rate of 30%, mitigated overfitting. The output layer utilized a sigmoid activation function to predict binary survival probabilities for each time interval (1-year to 5-year survivability), directly addressing our multi-temporal survival prediction objective.

3.4.2 Loss function and optimization

Loss Function: The ANN was trained using binary cross-entropy loss to produce calibrated probabilities suitable for survival prognosis and risk stratification (Appendix 5).

Optimization: Optimization employed the Adam algorithm, supporting stable learning under class imbalance and heterogeneous clinical covariates (Appendix 5).

3.4.3 Training and validation strategy

To explicitly prevent temporal leakage, we used a diagnosis-year split: 2010–2014 for training (N=8,943) and 2015 for testing (N=2,035). Within training, a stratified split created a validation subset for tuning. Training spanned 200 epochs specifically chosen to allow convergence for medical data complexity with a batch size of 64, balancing computational efficiency with gradient estimation stability.

3.5 Evaluation of neural network architectures

3.5.1 Overview of architectures

Various ANN architectures were systematically evaluated to identify the optimal configuration for oral cancer survival prediction: DenseNet-inspired, wide-and-shallow, deep-and-uniform, and symmetrical pyramidal designs.

3.5.2 Neural network transformations

Layer Transformations: Each layer applied a transformation:

$$z^{(l)} = W^{(l)}a^{(l-1)} + b^{(l)}, \quad a^{(l)} = f(z^{(l)}), \quad (1)$$

where $W^{(l)}$ and $b^{(l)}$ are the weight matrix and bias vector for layer l , and $f(x) = \ln(1 + e^x)$ is the Softplus activation function.

3.5.3 Regularization techniques

Batch Normalization: Batch normalization stabilized training and supported consistent risk stratification across heterogeneous patients (Appendix 5).

Dropout Regularization: Dropout at 0.3 reduced overfitting to site- or stage-specific patterns and improved out-of-sample prognosis (Appendix 5).

3.5.4 Architectural configurations

The following architectures were tested with specific rationales for oral cancer survival prediction:

- *DenseNet-inspired:* [128, 64, 32] This architecture uses a decreasing number of neurons per layer to progressively capture more abstract features, inspired by DenseNet principles to balance complexity and learning efficiency for medical data.
- *Wide and shallow:* [256, 128] Designed to capture broad feature representations with fewer layers, this structure aims to learn diverse clinical features quickly, beneficial when the dataset contains multiple heterogeneous clinical variables.
- *Deep and uniform:* [64, 64, 64, 64] Consists of multiple layers with equal width, promoting uniform feature depth across the network, which helps in learning hierarchical clinical relationships at multiple levels.
- *Symmetrical pyramid:* [32, 64, 128, 64, 32] Scaling the network width up and then down symmetrically attempts to balance feature expansion and compression, exploring both complexity and feature refinement, though it may not always be optimal for clinical datasets.
- *Transformer-like:* [128, 128, 128, 128] Inspired by transformer networks, this architecture maintains consistent width in layers to preserve feature dimension and allow complex clinical feature interactions crucial for survival prediction.

- *High capacity, deep*: [512, 256, 128] A deeper network with high initial capacity aims to capture complex clinical representations but risks overfitting without sufficient medical data, making regularization critical.
- *ResNet-inspired alternating*: [128, 64, 128, 64] Alternating layer widths reflect residual connections conceptually, aiming to improve gradient flow and allow the network to learn efficiently from clinical patterns.
- *Gradual reduction*: [256, 128, 64, 32] Starts with high dimensional input compression that gradually reduces complexity, allowing the network to abstract clinical features progressively while reducing parameters for better generalization in medical applications.

3.5.5 Experimental results

Each model was trained for 50 epochs with a batch size of 128. The transformer-like architecture [128, 128, 128, 128] achieved the strongest accuracy/F1 profile, while the gradual reduction architecture [256, 128, 64, 32] achieved the strongest AUC, indicating robust discrimination for oral cancer survivability prediction.

3.6 Observation

For oral cancer survivability prediction, transformer-like and gradual reduction architectures emerged as the most effective configurations, while the symmetrical pyramid design was comparatively weaker. This finding suggests that progressive feature abstraction is crucial for medical survival prediction tasks. Future work will focus on hyperparameter optimization and additional regularization techniques to enhance model performance.

3.6.1 Comparison of activation functions

We compared the performance of different activation functions on a fixed neural network architecture. Each activation function is defined as follows:

- *ReLU (Rectified Linear Unit)*:

$$f(x) = \max(0, x). \quad (2)$$

ReLU activates the neuron if the input is positive and outputs zero otherwise, making it computationally efficient as seen in Lin et al.(2018)[9].

- *Sigmoid*:

$$f(x) = \frac{1}{1 + e^{-x}}. \quad (3)$$

Sigmoid maps the input to a range between 0 and 1, making it suitable for probabilistic interpretations but prone to vanishing gradients as shown by Pratiwi et al.(2020) [10].

- *Tanh (Hyperbolic Tangent):*

$$f(x) = \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}. \quad (4)$$

Tanh maps the input to a range between -1 and 1, often leading to faster convergence compared to sigmoid as shown by Lau et al.(2018) [11].

- *ELU (Exponential Linear Unit):*

$$f(x) = \begin{cases} x & \text{if } x > 0, \\ \alpha(e^x - 1) & \text{if } x \leq 0, \end{cases} \quad (5)$$

where $\alpha > 0$ is a learnable parameter. ELU reduces bias shifts and helps with learning robust features as shown by Dubey et al.(2022) [12].

- *Softplus:*

$$f(x) = \ln(1 + e^x). \quad (6)$$

Softplus is a smooth approximation of ReLU, useful for situations where differentiability is important as shown by Nag et al.(2023) [13].

- *Leaky ReLU:*

$$f(x) = \begin{cases} x & \text{if } x > 0, \\ \alpha x & \text{if } x \leq 0, \end{cases} \quad (7)$$

where α is a small constant (e.g., 0.01). Leaky ReLU addresses the “dying ReLU” problem by allowing small gradients for negative inputs as shown by Xu et al.(2020) [14].

- *SELU (Scaled Exponential Linear Unit):*

$$f(x) = \begin{cases} \lambda x & \text{if } x > 0, \\ \lambda \alpha(e^x - 1) & \text{if } x \leq 0, \end{cases} \quad (8)$$

where $\lambda > 1$ and $\alpha > 0$ are fixed constants. SELU is designed to induce self-normalizing properties in the network as shown by Rasamoelina et al.(2020)[15].

The model architecture used in this experiment was three fully connected layers with 256, 128, and 64 neurons. Batch normalization and dropout layers followed the connected layers to further improve training. The impact of various activation functions

on the model is assessed using different evaluation metrics. These metrics were accuracy, F1 score, and AUC. This will provide clarity on how well the model performed under the various configurations.

3.6.2 Impact of dropout rates

To understand the effect of regularization on the model's generalizability, we experimented with varying dropout rates ($p = 0.2, 0.3, 0.4$, and 0.5) using a transformer-inspired architecture. The network consisted of four fully connected layers, each with 128 neurons, and utilized the `softplus` activation function. The dropout mechanism was applied to reduce overfitting and enhance the model's robustness.

Dropout and inference-time scaling are defined in Appendix 5; this subsection focuses on generalizability changes across dropout rates.

3.6.3 Transformer-inspired architecture with optimized dropout

In this experiment, we implemented a transformer-inspired neural network architecture for the task of survival prediction. The model architecture was designed with four fully connected layers, each comprising 128 neurons and using the `softplus` activation function. A dropout rate of 0.3, determined as optimal in prior experiments, was applied after each layer to improve generalizability. Batch normalization layers were also employed to stabilize and accelerate training.

Methodology:

- *Data Preprocessing:* The dataset was preprocessed by one-hot encoding categorical features and standardizing numerical features. These were then combined into a single feature matrix.
- *Data Splitting:* The dataset was partitioned into training (diagnosis years 2010–2014) and testing (diagnosis year 2015) sets to avoid temporal leakage; stratification was used only within training for validation.
- *Model Architecture:* The network comprised four dense layers, each structured as follows:

- `Dense(128, activation= softplus)`
- `BatchNormalization()`
- `Dropout(0.3)`

- *Output Layer:* The output layer employed a sigmoid activation function to perform binary classification.
- *Training Configuration:* The model was optimized using the Adam optimizer with a learning rate of 0.0003 and trained with a binary cross-entropy loss function. A batch size of 64 was utilized, and training was conducted over 200 epochs. To prevent overfitting and promote efficient convergence, a `ReduceLROnPlateau` callback, as described by Torghabeh et al. (2023) [16], was implemented to decrease the learning rate when validation loss stagnated.
- *Evaluation Metrics:* The model was evaluated using metrics such as accuracy

and F1 score and a detailed classification report. Additionally, trends in training and validation accuracy and loss were monitored during training, and the ROC curve was analyzed for model performance. Metric definitions are provided in Appendix 5; here we emphasize how these metrics support prognosis and risk stratification. *Significance:* The experiment demonstrated the effectiveness of transformer-inspired dense architectures, combining depth and regularization for robust survival prediction. The incorporation of dropout and batch normalization highlighted the importance of regularization in enhancing generalizability while mitigating overfitting.

Insights: The findings emphasized the importance of:

- Proper regularization techniques, such as dropout, to prevent overfitting.
- Stabilizing training through batch normalization.
- Optimizing architectural configurations to balance model depth and capacity with performance.

These insights contribute to the development of advanced architectures tailored for survival prediction tasks, leveraging deep learning methods effectively.

4 Result and discussion

4.1 Baseline dataset comparison

To establish the clinical significance of our neural network approach, we compared our model performance against established baseline methods commonly used in oral cancer survival prediction. Traditional approaches using Cox proportional hazards models and Random Survival Forest methods typically achieve accuracies of 60-65% and F1 scores of 65-70% on similar SEER datasets [1]. Our neural network models demonstrate consistent superiority across all configurations, with improvements ranging from 5.7% to 12.7% in accuracy and 8.9% to 13.9% in F1 scores, establishing the clinical value of deep learning approaches for oral cancer prognosis.

4.2 Comparison of activation functions

The results of the first experiment are summarized below. Table 4 highlights the performance of various activation functions in terms of test accuracy and F1 score.

Scientific Interpretation: The superior performance of Softplus (72.68% accuracy, 78.98% F1) over baseline methods (+7.7% accuracy improvement) demonstrates its ability to capture the complex, non-linear relationships inherent in oral cancer survival prediction. The smooth, differentiable nature of Softplus enables effective gradient propagation, crucial for learning intricate patterns in clinical data where survival outcomes depend on subtle interactions between patient demographics, tumor characteristics, and treatment modalities.

Table 4 Comparison of activation functions based on accuracy and F1 score

Activation function	Test accuracy (%)	F1 Score (%)
ReLU	71.72	78.91
Sigmoid	70.72	77.35
Tanh	71.22	77.81
ELU	72.27	78.70
Softplus	72.68	78.98
Leaky ReLU	72.22	78.94
SELU	72.40	78.93

Table 5 Comparison of layer configurations based on accuracy and F1 score

Layer configuration	Test accuracy (%)	F1 score (%)
[128, 64, 32]	72.45	79.30
[256, 128]	72.45	79.32
[64, 64, 64, 64]	71.99	78.83
[256, 128, 64, 32]	72.54	79.16
[128, 128, 128, 128]	72.59	79.26
[512, 256, 128]	72.18	78.55
[128, 64, 128, 64]	72.18	78.75
[32, 64, 128, 64, 32]	71.31	78.31

The poor performance of Sigmoid activation (70.72% accuracy) validates theoretical concerns about vanishing gradient problems in deep networks processing medical data. This finding has direct clinical implications: activation functions that fail to preserve gradient information may miss critical prognostic factors, potentially leading to suboptimal patient risk stratification. The 1.96% accuracy difference between Softplus and Sigmoid corresponds to approximately 215 additional correctly classified patients in our cohort of 10,978 cases (95% CI: 85–346 under a normal approximation for the accuracy difference).

4.3 Evaluation of layer configurations

The evaluation of different layer configurations revealed that the gradual reduction configuration ([256, 128, 64, 32]) achieved the highest AUC (0.76388), indicating its ability to balance feature abstraction and generalization effectively. The symmetrical pyramid configuration ([32, 64, 128, 64, 32]) had the lowest accuracy and F1 score, suggesting that overly symmetrical designs might not capture the complexity of the dataset adequately. (Table 5) Overall, configurations with gradual reductions and moderately deep architectures demonstrated robust performance across all metrics.

Clinical Significance of Architecture Choice: The superior performance of moderately deep configurations [128, 128, 128, 128] (72.59% accuracy) aligns with medical literature suggesting that oral cancer prognosis involves hierarchical feature relationships. Early layers learn basic clinical patterns (demographics, staging), while deeper layers capture complex interactions between treatment modalities and patient-specific factors. This architectural insight is clinically relevant as it mirrors the decision-making process oncologists use when considering multiple interdependent prognostic factors.

Table 6 Comparison of dropout rates based on test accuracy and F1 score

Dropout rate (p)	Test accuracy (%)	F1 score (%)
0.2	71.86	78.42
0.3	72.31	79.01
0.4	72.22	79.05
0.5	71.90	78.67

Table 7 Performance metrics for 1-Year survivability prediction

Stage	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Stage 1	90.07	90.74	98.89	94.64
Stage 2	86.82	87.50	98.32	92.61
Stage 3	88.33	89.04	98.60	93.56
Stage 4	85.52	86.33	98.36	91.79
Combined	87.86	88.58	98.40	93.20

The underperformance of symmetrical pyramid architecture [32, 64, 128, 64, 32] (71.31% accuracy) suggests that oral cancer survival prediction benefits from progressive feature abstraction rather than feature compression-expansion cycles. This finding implies that clinical data contains predominantly forward-flowing information patterns, where initial patient characteristics and tumor staging progressively determine treatment outcomes without requiring feature reconstruction. (Table 6)

4.4 Impact of dropout rates

The impact of varying dropout rates ($p = 0.2, 0.3, 0.4, 0.5$) on the performance of a transformer-inspired neural network was evaluated. The network architecture consisted of four fully connected layers with 128 neurons in each layer, `softplus` activation, and a varying dropout rate applied after each layer. The results in terms of test accuracy and F1 score are summarized in the table below.

Regularization Insights for Medical AI: The optimal dropout rate of 0.3 achieving 72.31% accuracy represents a critical balance for medical applications. This finding suggests that 30% feature dropout provides sufficient regularization to prevent overfitting to specific patient subgroups while retaining enough information to capture rare but clinically significant patterns. (Tables 7, 8, 9 10 and 11 In medical contexts, overly aggressive regularization ($p=0.5$) may inadvertently suppress important but infrequent clinical presentations, potentially missing critical cases.

4.5 Deep analysis of activation function performance (Table 4)

The 1.96% accuracy gap between the best (Softplus: 72.68%) and worst (Sigmoid: 70.72%) performing activation functions represents significant clinical impact. In our dataset of 10,978 patients, this translates to approximately 215 cases where activation function choice could determine accurate vs. inaccurate survival prediction. ELU's strong performance (72.27% accuracy) demonstrates the importance of handling negative activation values effectively—a crucial consideration when processing normalized clinical features that may contain negative values after standardization.

Table 8 Performance metrics for 2-Year survivability prediction

Stage	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Stage 1	84.69	85.47	97.71	91.20
Stage 2	78.57	79.22	95.53	86.56
Stage 3	80.42	81.65	96.18	88.33
Stage 4	80.00	82.17	93.81	87.58
Combined	80.23	81.13	95.35	87.38

Table 9 Performance metrics for 3-Year survivability prediction

Stage	Accuracy (%)	Precision (%)	Recall (%)	F1 score (%)
Stage 1	81.12	81.97	96.22	88.57
Stage 2	73.84	75.03	90.45	81.94
Stage 3	76.37	78.06	91.88	84.40
Stage 4	80.00	81.25	95.42	87.73
Combined	75.99	76.87	93.00	84.16

Table 10 Performance metrics for 4-Year survivability prediction

Stage	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Stage 1	78.92	79.32	95.57	86.65
Stage 2	71.81	72.53	87.11	79.21
Stage 3	74.60	75.73	89.60	82.13
Stage 4	81.38	83.33	92.23	87.54
Combined	74.07	75.14	91.13	82.58

Table 11 Performance metrics for 5-Year survivability prediction

Stage	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Stage 1	76.02	77.01	91.96	83.86
Stage 2	70.77	72.10	80.41	76.00
Stage 3	73.34	74.41	88.53	80.91
Stage 4	80.00	83.65	87.88	85.72
Combined	72.64	74.01	85.84	79.51

The clustering of top-performing functions (Softplus, SELU, ELU, Leaky ReLU) within a narrow performance band (72.22–72.68%) suggests that the key differentiator is not the specific mathematical formulation, but rather the presence of smooth gradients and effective handling of the dying neuron problem. This insight guides future model development: prioritizing activation functions that maintain gradient flow over those with specific biological interpretations.)

4.6 Model results for stage-wise output

This section presents the detailed performance metrics of the model across 1-year, 2-year, 3-year, 4-year, and 5-year survivability predictions. For each stage, the metrics include accuracy, precision, recall, and F1 score, expressed as percentages for the positive class. Additionally, the total combined metrics across all grades for each stage are reported as a separate row.

Clinical Interpretation of Temporal Survival Patterns: The systematic decline in model accuracy from 1-year (87.86%) to 5-year (72.64%) predictions reflects fundamental clinical realities of oral cancer progression. Early survival predictions benefit from immediately observable factors (tumor staging, initial treatment response), while long-term survival involves complex interactions with genetic factors, lifestyle changes, and potential recurrence patterns not fully captured in our feature set.

The counterintuitive finding that Stage 4 cancers show relatively stable performance across time periods (85.52% to 80.00% accuracy) versus Stage 2 cancers (86.82% to 70.77%) suggests that advanced-stage cancers follow more predictable progression patterns, albeit with poorer outcomes. This aligns with clinical observations that Stage 4 cancers, while having worse prognosis, demonstrate more consistent disease trajectories compared to intermediate stages where treatment response variability significantly impacts long-term outcomes.

The high recall values across all stages (>85%) indicate our model's strength in identifying patients at risk for mortality—a clinically crucial characteristic that minimizes false reassurances about survival prospects. However, the declining precision in longer-term predictions suggests increasing false positive rates, which in clinical practice could lead to unnecessarily aggressive treatment strategies for patients with better actual prognosis than predicted.

4.7 Decision-oriented thresholding and prediction-to-decision linkage

The models output a predicted probability of survival at horizon t , denoted $\hat{p}_i^{(t)} = \mathbb{P}(y_i^{(t)} = 1 \mid x_i)$. For operational decision-making, it is often more natural to work with predicted risk $r_i^{(t)} = 1 - \hat{p}_i^{(t)}$, where higher values indicate higher estimated risk of not surviving to horizon t . A basic linkage from prediction to action is thresholding: $a_i^{(t)} = \mathbb{I}\{r_i^{(t)} \geq \tau_t\}$, where $a_i^{(t)} = 1$ indicates a decision to allocate additional follow-up intensity (or referral) to patient i under horizon t . In practice, τ_t should not be chosen arbitrarily. A clinically meaningful reporting standard is recall-targeted thresholding: for a required sensitivity/recall level ρ (e.g., $\rho = 0.90$), choose the smallest threshold τ_t on the validation set such that $\text{Recall}(\tau_t) \geq \rho$, and then report the corresponding precision, the number of patients flagged, and the implied workload. This makes the deployment trade-off explicit: higher recall reduces missed high-risk patients but increases follow-up burden due to more false positives. When capacity constraints are explicit, the same linkage can be expressed as a constrained optimization: select a subset $\mathcal{S}_t$ of patients for additional follow-up to maximize expected utility subject to a capacity limit $|\mathcal{S}_t| \leq K_t$. A standard operational sur-

rogate is to rank patients by $r_i^{(t)}$ and choose the top K_t (equivalently, a threshold). A complete decision-analytic model would specify the utility/cost of true positives, false positives, true negatives, and false negatives, and would incorporate calibration checks to ensure portability across demographic subgroups. We do not implement and validate this full decision-analytic layer in the current scope; instead, we provide the above definitions and explicitly position formal decision optimization as prioritized future work (Sect. 5).

4.8 Operational interpretation: horizon choice, follow-up burden, and capacity trade-offs

Operationally, shorter horizons (1–2 years) can prioritize near-term surveillance and supportive care planning, while longer horizons (4–5 years) align with survivorship planning and resource forecasting. The decline in performance with longer horizons reflects the trade-off between lead time and uncertainty: in constrained settings, high-recall short-horizon screening can reduce missed high-risk cases, with longer-horizon predictions used primarily for patients with stable early follow-up trajectories. Stage-4 stability suggests more protocol-driven follow-up, while Stage-2 variability supports adaptive follow-up cadence where heterogeneity in response dominates long-term prognosis.

4.9 Fairness considerations and subgroup generalizability

Given the demographic skew noted in Section 3 (76.4% non-Hispanic White and approximately two-thirds male), fairness and generalizability require subgroup evaluation. In future work, we will report subgroup-stratified performance for race, sex, and age bands. Specifically, for each subgroup g we will report AUC_g , $F1_g$, $Recall_g$, and calibration error (e.g., expected calibration error) to assess whether predicted risks are comparably reliable across groups. We will also report disparity summaries such as the maximum absolute gap in recall across groups, $\max_{g,g'} |Recall_g - Recall_{g'}|$, and analogous gaps for precision and calibration. If disparities are observed, mitigation strategies (reweighting, subgroup-aware calibration, or constrained learning objectives) will be explored. In the current scope we report overall and stage-wise performance with temporal validation, and we treat subgroup fairness evaluation as a required next step for decision-relevant deployment.

5 Conclusion and future work

This research evaluates the impact of key design elements and training strategies to develop solutions for complex classification problems in cancer progression prediction. The study underscores the crucial role of artificial intelligence (AI) and mathematical modeling in addressing oral cancer challenges, which contribute significantly to cancer-related mortality. Neural networks offer a promising approach for improving cancer survivability predictions across 1-year to 5-year survivability for different oral cancer stages.

The declining accuracy from 87.86% (1-year) to 72.64% (5-year) reflects the inherent uncertainty in long-term prognosis, consistent with clinical observations. The superior performance in Stage I cancers demonstrates the model's ability to capture the biological reality that early-stage cancers have better prognostic outcomes, validating our model's clinical relevance.

The optimal dropout rate of 0.3 and effectiveness of Softplus activation function provide practical insights for developing robust neural networks in medical applications. The gradual decline in precision and recall over longer prediction periods suggests our model appropriately captures increasing prognostic uncertainty rather than providing false confidence in long-term predictions.

Despite limited datasets, this research has made significant strides in developing models that predict cancer progression and survival. The findings highlight AI's potential to analyze cancer stages, treatment modalities, and cancer spread, enhancing overall survivability of oral cancer patients. Even with limited information, AI-driven models can significantly improve diagnostic accuracy and support timely interventions.

Continued research with comprehensive datasets will further refine these models and bolster their effectiveness. Future investigations should incorporate baseline comparisons with traditional statistical methods such as Cox proportional hazards models to establish clinical superiority of neural network approaches. External validation using independent datasets from different geographic populations is essential to assess model generalizability.

The inclusion of behavioral risk factors such as smoking history and alcohol consumption in future models could significantly improve predictive accuracy. Furthermore, longitudinal analysis incorporating treatment response patterns and tumor biomarkers would enhance the model's ability to personalize treatment recommendations and optimize patient outcomes.

Finally, a formal prediction-to-decision framework is a natural next step. Future work will develop an explicit decision model mapping predicted survival risks to operational actions (e.g., follow-up cadence and referral thresholds), enabling quantification of clinical utility and policy impact beyond predictive accuracy alone.

In addition, future work will (i) benchmark horizon-based classification against time-to-event survival objectives (e.g., Cox/RSF/DeepSurv) on the same cohort, (ii) incorporate censoring-aware sensitivity analyses such as inverse probability of censoring weighting to quantify any remaining bias, and (iii) report subgroup-stratified performance and calibration (race, sex, age bands) to assess fairness and generalizability for deployment.

Appendix A: Mathematical Definitions

A.1 Binary cross-entropy loss

For binary survival classification, the training process minimized binary cross-entropy loss:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (9)$$

where y_i is the ground truth label, $\hat{y}_i$ is the predicted probability, and N is the number of samples.

A.2 Adam optimization update

Optimization used the Adam algorithm with dynamic learning rate adjustments:

$$\theta_t = \theta_{t-1} - \frac{\eta}{\sqrt{v_t} + \epsilon} m_t, \quad (10)$$

where m_t and v_t are the first and second moment estimates of the gradient as given by Bock et al.(2019)[17].

A.3 Batch normalization (BN)

Batch normalization was applied as:

$$\hat{x}_i = \frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}}, \quad y_i = \gamma \hat{x}_i + \beta, \quad (11)$$

where μ_B and σ_B^2 are the batch mean and variance, and γ, β are learnable parameters as shown by Santurkar et al.(2018) [18].

A.4 Dropout definition and inference-time scaling

For a neuron x_i , dropout with rate p during training is:

$$x_i^{\text{dropped}} = \begin{cases} 0 & \text{with probability } p, \\ x_i & \text{with probability } 1 - p. \end{cases} \quad (12)$$

During inference, outputs are scaled by $1 - p$:

$$x_i^{\text{scaled}} = x_i^{\text{dropped}} \cdot (1 - p). \quad (13)$$

A.5 Evaluation metric definitions

Accuracy:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}. \quad (14)$$

F1 score:

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (15)$$

where Precision = $\frac{TP}{TP+FP}$ and Recall = $\frac{TP}{TP+FN}$.

Author contributions Raghav Aggarwal conceptualized the research, conducted the analysis, and authored the manuscript. S Dharmaraja provided mentorship, technical guidance, and critical feedback throughout the research and manuscript preparation process. Both authors reviewed and approved the final manuscript.

Funding This is a non funded retrospective study.

Data availability SEER Research Data, 12 Registries, Nov 2023 Sub (1992-2021) - Linked To County Attributes - Time Dependent (1990-2022) Income/Rurality, 1969-2022 Counties analyzed for this study can be found at the following link: <https://seer.cancer.gov/> SEER Research Data 1975-2021 when Using SEER*Stat:

Declarations

Conflict of interest The authors declare that they have no Conflict of interest.

Ethical approval Manuscripts reporting studies involving human participants, human data or human tissue: Not applicable

Consent for publication This manuscript does not contain data from any individual person, hence, it is “Not applicable.”

References

1. Kim, D.W., Lee, S., Kwon, S., Nam, W., Cha, I.-H., Kim, H.J.: Deep learning-based survival prediction of oral cancer patients. *Sci. Rep.* **9**(1), 6994 (2019)
2. Bray, F., Laversanne, M., Sung, H., Ferlay, J., Siegel, R.L., Soerjomataram, I., Jemal, A.: Global cancer statistics 2022: globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A Cancer J. Clin.* **74**(3), 229–263 (2024)
3. Tabassum, S., Rosli, N.B., Binti Mazalan, M.S.A.: Mathematical modeling of cancer growth process: a review. In: *Journal of Physics: Conference Series*, vol. 1366, p. 012018 (2019). IOP Publishing
4. Teng, J., Zhang, H., Liu, W., Shu, X.-O., Ye, F.: A dynamic bayesian model for breast cancer survival prediction. *IEEE J. Biomed. Health Inform.* **26**(11), 5716–5727 (2022)
5. Gang, Q., Feng, J., Kauczor, H.-U., Zhang, K.: Predicting nodal metastasis progression of oral tongue cancer using a hidden markov model in mri. *Front. Oncol.* **14**, 1360253 (2024)
6. Rajaguru, H., Prabhakar, S.K.: Performance comparison of oral cancer classification with gaussian mixture measures and multi layer perceptron. In: *The 16th International Conference on Biomedical Engineering: ICBME 2016, 7th to 10th December 2016, Singapore*, pp. 123–129 (2017). Springer
7. Weatherspoon, D.J., Chattopadhyay, A., Boroumand, S., Garcia, I.: Oral cavity and oropharyngeal cancer incidence trends and disparities in the united states: 2000–2010. *Cancer Epidemiol.* **39**(4), 497–504 (2015)
8. Baik, J., Ye, Q., Zhang, L., Poh, C., Rosin, M., MacAulay, C., Guillaud, M.: Automated classification of oral premalignant lesions using image cytometry and random forests-based algorithms. *Cell. Oncol.* **37**, 193–202 (2014)

9. Lin, G., Shen, W.: Research on convolutional neural network based on improved relu piecewise activation function. *Procedia Comput. Sci.* **131**, 977–984 (2018)
10. Pratiwi, H., Windarto, A.P., Susliansyah, S., Aria, R.R., Susilowati, S., Rahayu, L.K., Fitriani, Y., Merdekawati, A., Rahadjeng, I.R.: Sigmoid activation function in selecting the best model of artificial neural networks. In: *Journal of Physics: Conference Series*, vol. 1471, p. 012010 (2020). IOP Publishing
11. Lau, M.M., Lim, K.H.: Review of adaptive activation function in deep neural network. In: *2018 IEEE-EMBS Conference on Biomedical Engineering and Sciences (IECBES)*, pp. 686–690 (2018). IEEE
12. Dubey, S.R., Singh, S.K., Chaudhuri, B.B.: Activation functions in deep learning: a comprehensive survey and benchmark. *Neurocomputing* **503**, 92–108 (2022)
13. Nag, S., Bhattacharyya, M., Mukherjee, A., Kundu, R.: Serf: Towards better training of deep neural networks using log-softplus error activation function. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 5324–5333 (2023)
14. Xu, J., Li, Z., Du, B., Zhang, M., Liu, J.: Reluplex made more practical: leaky relu. In: *2020 IEEE Symposium on Computers and Communications (ISCC)*, pp. 1–7 (2020). IEEE
15. Rasamoelina, A.D., Adjailia, F., Sinčák, P.: A review of activation function for artificial neural network. In: *2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMI)*, pp. 281–286 (2020). IEEE
16. Torghabeh, F.A., Modaresnia, Y., Khalilzadeh, M.M.: Effectiveness of learning rate in dementia severity prediction using vgg16. *Biomed. Eng.: Appl. Basis Commun.* **35**(03), 2350006 (2023)
17. Bock, S., Weiß, M.: A proof of local convergence for the adam optimizer. In: *2019 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8 (2019). IEEE
18. Santurkar, S., Tsipras, D., Ilyas, A., Madry, A.: How does batch normalization help optimization? *Adv. Neural. Inf. Process. Syst.* **31** (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.