VET nuus•news

Oktober / October 2025

The Monthly Magazine of the SOUTH AFRICAN VETERINARY ASSOCIATION
Die Maandblad van die SUID-AFRIKAANSE VETERINÊRE VERENIGING

THEME

Opthalmology

CPD

Reproductive disorders of backyard poultry

South African Veterinary Association
Suid-Afrikaanse Veterinêre Vereniging

Access to CPD Articles:

CPD CONTINUING PROFESSIONAL DEVELOPMENT

https://www.sava.co.za/vetnews-2025/

Dagboek • Diary

Ongoing / Online 2025

SAVETCON: Webinars
Info: Corné Engelbrecht, SAVETCON, 071 587 2950,
corne@savetcon.co.za / https://app.livestorm.co/svtsos

Acupuncture – Certified Mixed Species Course
Info: Chi University:
https://chiu.edu/courses/cva#aboutsouthafrica@tcvm.com

SAVA Johannesburg Branch CPD Events
Monthly - please visit the website for more info.
Venue: Johannesburg Country Club
Info: Vetlink - https://savaevents.co.za/

October 2025

Oranje Vaal CPD Day
11 October
Venue: Afridome, Parys
Info: conference@savetcon.co.za

Northern Natal and Midlands Branch Congress
11-12 October
Venue: Fordoun Hotel and Spa, Midlands
Info: https://vetlink.co.za/northern_natal_and_midlands/

The Middle East & Africa Veterinary Congress (MEAVC)
17 -19 October: Pre- and Main Congress Workshops
Venue: Jafza One Convention Centre, Dubai
Info: www.meavc.com

SAVA Free State and Northern Cape Branch Congress
17-18 October
Venue: Goose Hill Guest Farm, Bloemfontein
Info: conference@savetcon.co.za

KwaZulu-Natal Branch Congress
25-26 October
Venue: San Lameer Resort, Southbroom
Info: www.vetlink.co.za

11th International Sheep Veterinary Congress
27-31 October
Venue: Wollongong, New South Wales, Australia
Info: https://intsheepvetassoc.org/11th-isvc-2025

10th Annual South African Immunology Society (SAIS) Congress
30 October – 01 November
Venue: Garden Court Marine Parade, Durban (KZN)
Info: corne@savetcon.co.za or visit www.savetcon.co.za

Southern Cape Branch Congress
31 October – 01 November
Venue: Oubaai Hotel Golf & Spa, George
Info: https://vetlink.co.za/southern-cape-branch/

February 2026

SAEVA Congress 2026
19 - 22 February
Venue: Champagne Sports Resort, Drakensberg, KwaZulu Natal
Info: https://vetlink.co.za/saeva-congress-2026/

March 2026

Wildlife Group of the SAVA Congress 2026
12 - 14 March
Venue: 26° South Hotel , Muldersdrift, Gauteng
Info: https://vetlink.co.za/wildlife-group-2026/

NVCG: The Veterinary Masterclass: Neurology
17 - 18 March
Venue: Cape Town (venue to be confirmed)
19 - 20 March
Venue: Johannesburg (venue to be confirmed)
Info: Vetlink: www.nvcg.co.za or +2712 346 1590

June 2026

13th Int Crustacean Society Mid-Year Meeting
01-04 June
Venue: STIAS – Stellenbosch
Info: https://tcs2026.com/

RuVASA Congress 2026
08 - 10 June
Venue: Champagne Sports Resort, Drakensberg, KZN
Info: https://vetlink.co.za/ruvasa-congress-2026/

Septembr 2026

NVCG: The Vet Masterclass: Medicine & Dermatology
19 - 20 September
Venue: Johannesburg (venue to be confirmed)
21 - 22 September
Venue: Cape Town (venue to be confirmed)
Info: Vetlink: www.nvcg.co.za or +2712 346 1590

SAAVT Biennial Congress
30 September - 01 October
Venue: Krystal Beach Hotel – Gordon's Bay
Info: conference@savetcon.co.za

SAVE THE DATE

VET nuus•news

Contents I Inhoud

STREET ADDRESS
47 Gemsbok Ave, Monument Park, Pretoria, 0181, South Africa

POSTAL ADDRESS
P O Box 25033, Monument Park
Pretoria, 0105, South Africa

TELEPHONE
+27 (0)12 346-1150

FAX
General: +27 (0) 86 683 1839
Accounts: +27 (0) 86 509 2015

WEB
www.sava.co.za

CHANGE OF ADDRESS
Please notify the SAVA by email:
debbie@sava.co.za or letter: SAVA,
P O Box 25033, Monument Park, Pretoria, 0105, South Africa

CLASSIFIED ADVERTISEMENTS
(Text to a maximum of 80 words)
Sonja van Rooyen
assistant@sava.co.za
+27 (0)12 346 1150

DISPLAY ADVERTISEMENTS
Sonja van Rooyen
assistant@sava.co.za
+27 (0)12 346 1150

DESIGN AND LAYOUT
Sonja van Rooyen

PRINTED BY
Business Print: +27 (0)12 843 7638

Diary / Dagboek

Regulars / Gereeld

Articles / Artikels

Association / Vereniging

Vet's Health / Gesondheid

Technical / Tegnies

Marketplace / Markplein

Jobs / Poste

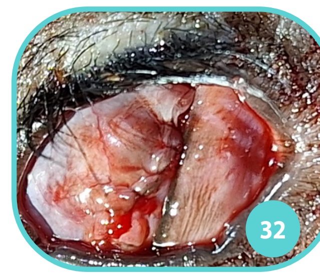

24-Hour Toll-Free Helpline: 0800 21 21 21

SAVA
South African Veterinary Association
Suid-Afrikaanse Veterinêre Vereniging

CPD CONTINUING PROFESSIONAL DEVELOPMENT

Click on the image to access Vetnews CPD articles

President: Dr Ziyanda Qwalela president@sava.co.za
Interim General Manager: Ms Sonja Ludik sonja@sava.co.za/ +27 (0)12 346 1150
Editor VetNews: Ms Andriette van der Merwe vetnews@sava.co.za
Accounts / Bookkeeping: Ms Shaye Hughes accounts@sava.co.za/+27 (0)12 346 1150
Reception: Ms Hanlie Swart reception@sava.co.za/ +27 (0)12 346 1150
Marketing & Communications: Ms Sonja van Rooyen marketing@sava.co.za/ +27 (0)12 346 1150
Membership Enquiries: Ms Debbie Breeze debbie@sava.co.za/ +27 (0)12 346 1150
Vaccination Booklets: Ms Debbie Breeze debbie@sava.co.za/ +27 (0)12 346 1150
South African Veterinary Foundation: Ms Debbie Breeze savf@sava.co.za/ +27 (0)12 346 1150
Community Veterinary Clinics: Ms Claudia Cloete manager@savacvc.co.za/ +27 (0)63 110 7559
SAVETCON: Ms Corné Engelbrecht corne@savetcon.co.za/ +27 (0)71 587 2950

Dear members,

Let us keep the pace and remain impactful!

What a whirlwind month September has been! I would like to congratulate and thank all our colleagues who freely gave of their time and skills to support the vaccination, awareness, and sterilisation campaigns conducted during this period. Your generosity does not go unnoticed, and it reflects the profession's enduring commitment to animal health, one health and the public good.

The constitution of the South African Veterinary Council (SAVC) continues to dominate conversations within the fraternity. At its meetings held on 26 and 27 September respectively, the SAVA Board of Directors and Federal Council considered the ongoing impasse regarding the constitution of the SAVC, the notice of motion received as part of the legal proceedings between the SAVC and the Minister of Agriculture, as well as various communications from members and Special Interest Groups.

Both meetings resolved to create a poll to allow members to express their preferences regarding the matter. Several possible approaches have been identified to address the situation. These options are not mutually exclusive, and any course of action will require cooperation by the parties concerned. This poll is intended solely to gauge the preferences of SAVA members. While the aggregated results will be shared with the relevant parties where applicable, the preferences expressed will not necessarily determine or influence the outcome of the matter. It was also clarified at FEDCO that despite some written communication, the parties have not yet formally met to resolve the issue. Both meetings emphasised that the desired end state is the constitution of a credible veterinary council within the legal framework as soon as possible. The President has accordingly been tasked to engage the National Department of Agriculture and Ministry to create a space for dialogue and the resolution of the impasse.

In light of these discussions, the meetings also reflected on how we as colleagues use SAVA platforms to engage on important matters. While these platforms cannot be regulated completely, there will always be a need for a degree of self-awareness and discernment about what and how we post. We are all passionate about the profession, and that passion is a strength. At the same time, let us be mindful of how we express ourselves: engaging in good faith, with integrity, and with mutual respect. This approach ensures that robust debates remain valuable learning opportunities for all and strengthen, rather than fracture, our professional community. This consciousness is amplified when it comes to senior members of the profession as well as portfolio holders of the organisation.

On behalf of the Board and Federal Council, I would like to extend our sincere appreciation to our outgoing SAVA representative on the SAVC, Dr Brendan Tindall. Your dedication, professionalism, and tireless advocacy for the interests of the profession have been invaluable. We are deeply grateful for the time, expertise, and commitment you have given to uphold high standards of governance and ethical practice, and we wish you well in your future endeavours.

We also wish to recognise the contribution of our incoming SAVA representative on the SAVC, Dr. Leon De Bruyn. Stepping into this role has been something of a baptism of fire, yet you have already contributed significantly to creating balance and perspective in the ongoing debates around the SAVC. Your willingness to engage constructively at such a critical time is commendable, and we look forward to your continued leadership and insight as we navigate these challenges together.

Similarly, our new Interim General Manager, Sonja Ludik, has also experienced something of a baptism of fire and is already thriving in her role. We warmly welcome you, Sonja, and trust that you will find the challenges both stimulating and rewarding, and that the experience gained will prove invaluable in your professional growth.

Another issue discussed was the live export by sea of animals destined for slaughter. Both meetings agreed that a similar poll of members should be conducted on this subject. The SAVA Animal Ethics and Welfare Group with the assistance of Professor. Gareth Bath, who has contributed significantly in this area, has been tasked with drafting an appropriate statement for members to consider via the poll.

This issue is complex and calls on veterinarians to provide leadership grounded in ethical and scientific principles and practices, a clear understanding of the applicable legal framework, and awareness of the socio-economic impacts. It was highlighted at both meetings that SAVA can advise and provide guidance, but it is not responsible for regulatory functions and cannot dictate the implementation of legal instruments in this matter.

Coming up this month is the much-anticipated WVA VPP Workshop, to be held in Pretoria. The workshop will address concerns regarding the implementation of an FAO project assisting VPPs to open their own practices in three African countries, and will also explore the possibility of similar projects for veterinarians.

Nominations for attendance have been requested from the NVCG, RUVASA, and the branches where the project was piloted—namely the North West, Eastern Cape, and Karoo.

The Strategy Development Committee, working closely with FINCO, has continued to focus on safeguarding SAVA's financial viability. This has proved to be an uphill task given delayed HWSETA payments, the slow realisation of sponsorships, and outstanding membership fees. The key challenge remains the need for SAVA to operate strictly within its means, necessitating a leaner and more efficient administration. Despite this, SAVA has managed to reduced its operational costs significantly, a phenomenal performance and commendable achievement by all involved.

The issue of outstanding membership fees remains a critical concern. Branches and groups are therefore urged to actively encourage their members to regularise their subscriptions, as this revenue is essential to sustaining the work of the Association and ensuring its long-term effectiveness.

Thank you once again to all members who continue to give of their time, energy, and expertise. These collective efforts strengthen our profession and ensure that SAVA remains a credible voice for veterinarians. V

Groetnis!
Ziyanda
President of the South African Veterinary Association

Editor's notes / Redakteurs notas

A person goes from not being able to see, or to see with difficulty, to being able to see clearly. Something a lot of us take for granted. With World Eyesight Day this month, I want to encourage all to also you to look after your own eyes. Your eyes give you the ability to do your daily job and to enjoy the wonderful world around you. Do not neglect this one sense that brings so much joy.

Last month, we celebrated International Rabies Day, and I hope everybody has recovered from the myriad of activities that took place. In Hoedspruit, we had the sad case of a dog who missed out on a rabies campaign earlier and was subsequently diagnosed with rabies. A stark reminder that it is still out there, and every animal not vaccinated is at risk. But also puts the people at risk.

Vetnews thanks every person for every rabies shot given, every person educated in the importance of rabies control. We still strive for zero rabies in 2030.

We think of the hundreds of hectares of veld that were destroyed in Etosha and the surrounding farms. It is devastating after the good rains they had last summer. This map is a satellite image of 26 September 2025 and indicates fires in Southern Africa.

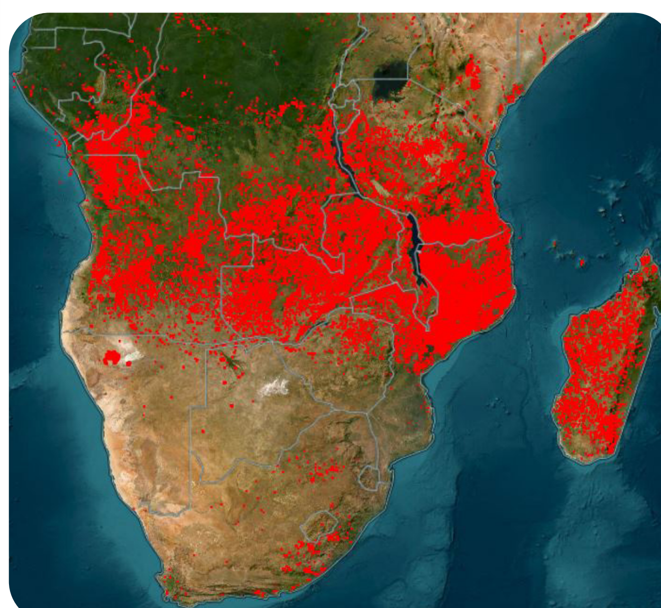

A scary picture indeed, but a good reflection of the seasonal fires of subsistence farmers that prepare for the new planting season. We pray for good rain in all the affected areas.

May your October be a colourful array of blessings.

Andriette

What is the eye?

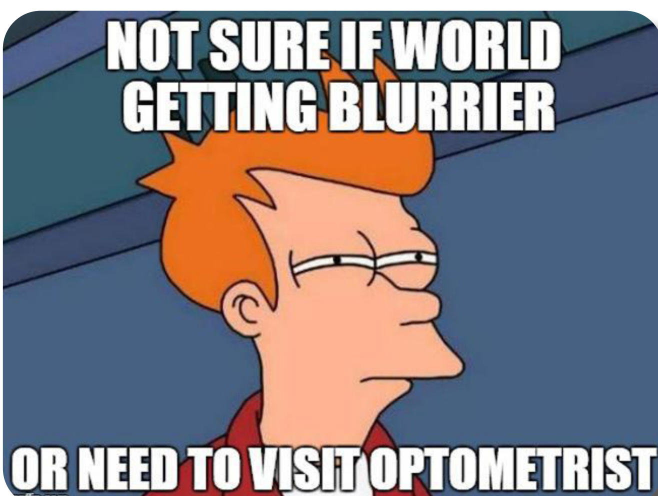

The eye is a fluid-filled bulb that allows us to observe by sight. A miracle in itself. At exactly the right pressure, it allows us to see. It is the complex sensory organ responsible for sight, working by focusing light onto the light-sensitive retina at the back of the eye.

The retina converts this light into electrical signals that the optic nerve transmits to the brain, creating an image we perceive. The eye's structures, including the protective outer layers, the vascular middle layer, and the light-sensitive retina, work together with fluids like aqueous and vitreous humour, and the focusing lens to achieve vision.

I watched a video of a lens replacement once, and I was mesmerised by the intricacy of eye surgery. To be able to do surgery on such a sensitive organ is just mind-blowing. My one son-in-law is stepping into the world of ophthalmology, and I am always stunned by the effect of cataract removal or other corrective surgeries.

« BACK TO CONTENTS

Hill's
PRESCRIPTION DIET
itching eased
joy retrieved
Derm/Helicopter/VetNews
NEW
PUPPY
FORMULA
Hill's
PRESCRIPTION DIET
DERM COMPLETE
skin care & food sensitivities
VETERINARIAN RECOMMENDED
PUPPY
Derm Complete Adult is also available in a Mini kibble for small dogs
Long-term support for their skin — and your treatment plan.
1
Clinically proven nutrition for both food and environmental allergies supported by 4 clinical studies
2
Nutrition formulated to support the skin barrier against environmental allergies – year round
SCIENCE DID THAT.
© Hill's Pet Nutrition, Inc. 2024
V32800; V32802; V35148

Save the Date

National Veterinary Clinicians Group of the SAVA
Advancing Veterinary Practice for Small Animal Vets

THE VETERINARY MASTERCLASS 2026 | MARCH NEUROLOGY

Stay ahead with cutting-edge expertise

17-18 March 2026
Cape Town

19 - 20 March 2026
Johannesburg

Steven De Decker
DVM PhD DipECVN MvetMed PGCertVetEd FHEA MRCVS

Steven De Decker graduated from Ghent University in Belgium. After graduation, he performed a rotating internship there and undertook a PhD studying 'wobbler syndrome' in dogs. This was followed by a Residency in Neurology and Neurosurgery at the Royal Veterinary College. He is Senior Lecturer and the Head of Service of the neurology and neurosurgery team at the Royal Veterinary College. Although he is interested in all aspects of veterinary neurology, most of his research and publications focus on spinal disorders and neurosurgery.

National Veterinary Clinicians Group of the SAVA
Advancing Veterinary Practice for Small Animal Vets

THE VETERINARY MASTERCLASS 2026 | SEPTEMBER MEDICINE & DERMATOLOGY

Stay ahead with cutting-edge expertise

19th - 20th September 2026
Johannesburg

21st - 22nd September 2026
Cape Town

SPEAKERS

Prof Ralf S Meuller
Prof. Dr.med.vet. Ralf S. Mueller, DipACVD, FANZCVSc (Dermatology), DipECVD

Prof Katrin Hartmann
Prof. Dr. med. vet., Dr. habil. Hartmann, Katrin, Prof., Dipl. ECVIM-CA (Internal Medicine)

www.nvcg.co.za

+2712 346 1590

Application of Computer Vision Methods in Veterinary Ophthalmology

An excerpt from a DOCTORAL THESIS by Matija Burić of the
UNIVERSITY OF RIJEKA FACULTY OF INFORMATICS AND DIGITAL TECHNOLOGIES

ABSTRACT

This dissertation applies computer vision techniques, leveraging deep learning models such as U-Net and GPT-4o, to improve the diagnosis of canine eye diseases in veterinary ophthalmology. The DogEyeSeg4 dataset of real-world clinical images serves as the foundation for training. Synthetic images augmented the dataset to enhance model robustness and generalisation.

U-Net(RSD), trained on DogEyeSeg4 and synthetic images generated using Stable Diffusion, was used for precise segmentation of canine eye symptoms such as corneal cloudiness, scleral redness, excessive tearing, and colored mass protrusion in the eye corner. The study also trained individual binary segmentation models for each symptom, utilising heatmaps from SSD eye detection to reduce false positives. Although these binary models improved symptom isolation, they faced challenges with overlapping conditions and increased complexity. Ultimately, the multiclass UNet(RSD) model provided better overall performance and efficiency.

GPT-4o interpreted the segmented images, outperforming other Large Language Models (LLMs) in generating accurate diagnostic suggestions, particularly when using segmentation masks from the adjusted U-Net with a ResNet backbone alongside the original images.

Despite promising results, challenges remain in diagnosing complex or subtle conditions like corneal ulcers. Future work includes expanding the dataset and symptom range, improving model architectures, and integrating multimodal data for more holistic diagnostics. These findings underscore the potential for AI-driven tools to revolutionise veterinary ophthalmology, offering more accurate and efficient diagnostic processes that can ultimately improve animal care.

1. INTRODUCTION

1.1. Problem and Research Subject with Hypothesis

The field of veterinary ophthalmology faces significant challenges in diagnosing and managing ocular conditions in canines, primarily due to the limited availability of advanced diagnostic tools. While human medicine has seen remarkable progress with the integration of computer vision techniques, especially in the domain of ophthalmology, these technological advancements have not been fully leveraged in veterinary applications. The disparity between the diagnostic capabilities available in human and veterinary medicine highlights a critical gap that this research aims to address. The core problem of this dissertation is to bridge this gap by developing and applying computer vision methods specifically tailored for canine eye disease detection and diagnosis.

The subject of this research centres on the innovative application of deep learning models, particularly those designed for image segmentation, to enhance the accuracy and efficiency of diagnosing ocular conditions in dogs. The research explores the potential of adapting and optimising these models, which have proven effective in human medical imaging, to meet the unique challenges posed by veterinary ophthalmology. This includes dealing with diverse imaging conditions, varying image quality, and the need for precise segmentation of specific ocular features.

To guide this research, the following hypotheses have been formulated:

- *A Computer vision model can recognise certain canine ocular conditions in still images taken in an unconstrained environment.*
- *Modification of the input and architecture of the U-Net network contributes to a better segmentation of canine eye conditions.*

The first hypothesis addresses the core objective of demonstrating that a computer vision model, when applied to images captured in non-ideal, real-world settings, can accurately identify specific ocular conditions in canines. This aspect of the research is crucial because, unlike controlled laboratory environments, real-world veterinary practices often involve images that vary significantly in quality due to factors such as lighting, angle, and the behaviour of the animals during imaging.

The second hypothesis related to the technical enhancement of the model itself, positing that by modifying both the input features and the architectural elements of the U-Net network, the model's performance in segmenting canine eye conditions can be improved. This hypothesis suggests that through careful adjustments and refinements, the model can be made more robust and capable of handling the complexities associated with veterinary ophthalmic images, which often include a wide range of conditions and variations.

Together, these hypotheses set the foundation for a comprehensive investigation into the applicability of advanced computer vision techniques in veterinary medicine, aiming to develop tools that could revolutionise the way canine eye diseases are diagnosed and managed in clinical settings.

1.2. Purpose and Objectives

The purpose of this dissertation is to apply advanced computer vision techniques to improve the diagnosis of canine eye diseases. The research focuses on developing a novel dataset and deep learning models for disease recognition and segmentation in real-world conditions. The specific objectives, aligned with the expected scientific contributions, are:

1. *An image dataset for machine learning of canine eye diseases:* Creation of a publicly available, annotated dataset for training machine learning models on canine eye diseases.
2. *Deep convolutional neural network model for recognition of canine eye clinical symptoms and diseases from still images in unconstrained environments:* Development of a CNN-based model for identifying eye symptoms and diseases in images captured under varied conditions.
3. *Deep neural network based on U-Net for segmentation of canine eye clinical symptoms from still images in unconstrained environments:* Implementation of a U-Net-based model for segmenting clinical symptoms in canine eye images.
4. *An improved method for segmentation of canine eye conditions based on U-Net:* Refinement of the U-Net model to improve segmentation accuracy for diagnosing canine eye conditions.

1.3. Brief Review of Previous Research

Computer vision has made significant strides in human medical diagnostics, with successful applications in ophthalmology, such as the detection of glaucoma and retinal diseases using Convolutional Neural Networks (CNNs).

However, the application of similar technologies in veterinary medicine, particularly for diagnosing canine eye diseases, is limited. Previous studies have demonstrated the potential of CNNs for disease recognition in controlled environments, but there remains a gap in their application under real-world, unconstrained conditions. This research aims to bridge that gap by developing a tailored dataset and a specialised U-Net model for veterinary use.

1.4. Scientific Methods

The research employs several scientific methods:

1. *Data Collection and Annotation*: A custom dataset was created using images collected from veterinary clinics and annotated by experts.
2. *Deep Learning Model Development:* The U-Net architecture was modified and trained using transfer learning and data augmentation techniques to improve segmentation accuracy.
3. *Model Evaluation:* The performance of the model was assessed using metrics such as the Jaccard Index and Dice Similarity Coefficient, followed by statistical analysis using ANOVA and Tukey HSD tests.
4. *Application Development:* A web-based tool was developed and deployed using Docker to ensure portability and scalability.

2. LITERATURE REVIEW

The application of computer vision techniques has significantly advanced the field of medicine, particularly in the diagnosis and analysis of various medical conditions. Deep **Convolutional Neural Networks** (CNNs), a class of deep learning models, have been widely used in the field for tasks such as detecting glaucoma, diabetic retinopathy, and other retinal disease recognition tasks, providing high accuracy in ocular image analysis and disease classification [1], [2]. These advancements are attributed to CNNs' ability to automatically extract and learn relevant features from ocular images, improving diagnostic accuracy compared to traditional manual methods [3], [4], [5].

The development of these CNN-based models has been supported by a variety of publicly available datasets, which have enabled the training and validation of models in human ophthalmology. Notable datasets include: ORIGA-light dataset focusing on optic nerve head segmentation and glaucoma assessment [6], Drishti-GS, designed for glaucoma detection, containing annotated fundus images for training CNNs [7], Retinal

Fundus Image for Glaucoma Detection with a focus on glaucoma diagnosis [8], RIMONE, a large dataset for retinal image analysis, primarily used for glaucoma detection [9], iBUG focuses on facial landmark detection but has also been used for eye-tracking and ocular disease analysis [10], OpenEDS dataset that includes annotated eye images for eye-tracking and ocular disease applications [11], UBIRIS primarily for biometric purposes, but also applied in ocular disease diagnosis [12] and TEyeD dataset containing eye-tracking data, useful for studying eye movements and diagnosing conditions such as glaucoma [13]. These datasets have significantly contributed to the development of effective models for human ophthalmology. In contrast, the application of computer vision techniques in veterinary ophthalmology, particularly for canine eye diseases, is much less developed. Datasets and research in this field are scarce. Studies focusing on conditions such as canine glaucoma are limited [14], [15], [16], and the few available datasets are typically small and lack the diversity seen in human datasets. A notable study used CNNs to diagnose ulcerative keratitis in dogs, but it was constrained by a limited dataset and acquired under controlled conditions [17]. This highlights the need for more extensive datasets and research into canine ophthalmology.

The U-Net architecture, introduced for biomedical image segmentation, has become a popular choice for various medical imaging tasks due to its encoder-decoder structure, which enables precise feature extraction and localisation [18]. U-Net's architecture is particularly effective in segmenting medical images where pixel-level accuracy is critical. It has demonstrated robust performance across multiple domains, including ophthalmology, where it has been applied to tasks such as optic disc segmentation, retinal layer segmentation, and detecting diseases like glaucoma and cataracts [19], [20], [21], [22]. Studies have consistently shown that U-Net performs well in medical image analysis, confirming its relevance in both human and veterinary medical contexts [18], [19], [20], [23], [24], [25]. Transformer-based models, such as Swin Transformer [26] and SegFormer [27], have recently gained attention for their superior ability to capture global context, making them particularly effective in fine-grained segmentation tasks across human and veterinary fields. However, U-Net continues to be a formidable option, especially in situations with limited data availability, where it still performs robustly [28]. In fact, a comparison between U-Net and transformer-based architectures for medical image registration demonstrated that UNet, with minimal adjustments, can surpass the performance of these newer transformer models [28].

One of the key strengths of U-Net is its adaptability, particularly when combined with transfer learning. Transfer learning allows models to leverage pre-trained weights from large datasets, enabling them to perform well even on smaller datasets typical in veterinary applications. U-Net has been successfully combined with advanced CNN backbones like ResNet [29], EfficientNet [30], VGG [31], and Inception [32], significantly improving its feature extraction capabilities [4], [5]. These backbones help enhance the model's ability to learn from limited data, making it well-suited for applications where obtaining large datasets is impractical. Furthermore, studies have shown that U-Net does not require an extensive dataset to achieve good results, especially when augmented with transfer learning techniques [33].

While U-Net can perform well with small datasets, the creation of synthetic datasets can further enhance model performance by augmenting the available training data. Synthetic data generation techniques, such as those using Generative Adversarial Networks (GANs) and diffusion models, have become valuable tools for augmenting datasets in various domains [34], [35]. Diffusion models, particularly Stable Diffusion, have shown promise in generating high-quality synthetic images by iteratively refining noisy images [36]. Stable Diffusion, which employs a U-Net-like architecture for image synthesis, can be used to create realistic synthetic images of diseased eyes, which can significantly improve the training of CNNs for canine ophthalmology [37], [38], [39]. These synthetic datasets could address the shortage of annotated data in veterinary ophthalmology and enhance model robustness [40].

A web-based application utilising a U-Net model trained on real and synthetic datasets could provide a valuable diagnostic tool for veterinarians. This application could assist in the early detection and treatment of canine eye diseases. In combination with image segmentation, the integration of Large Language Models (LLMs) could further enhance the diagnostic process. LLMs such as ChatGPT [41], Mistral [42], Gemini [43], Llama [44], and Claude [45] have demonstrated significant potential in medical data interpretation. These models, when integrated with image analysis tools, could help in interpreting symptoms, guiding veterinarians through diagnostic workflows, and improving decision-making. Evaluation metrics like BERTScore [46], CLIPScore [47], BLEU [48], METEOR [49], ROUGE [50], and SPICE [51] have shown that LLMs can effectively process complex medical information, making them suitable for integration into diagnostic applications.

This research focuses on developing a mobile web application for diagnosing canine eye diseases by combining U-Net-based image segmentation with LLMs for interpreting medical symptoms.

The U-Net model is trained on a custom dataset – DogEyeSeg4 [52], augmented with synthetic data generated using Stable Diffusion techniques. This study explores the integration of LLMs with medical image analysis, evaluating the potential of various LLMs such as ChatGPT, Mistral, Gemini, and others in improving diagnostic workflows. By combining these state-of-the-art technologies, this research aims to advance the field of veterinary ophthalmology, providing tools that can assist in the early detection and treatment of eye diseases in dogs.

3. METHODOLOGY

3.1. Dataset

In veterinary ophthalmology, especially in the niche of canine eye diseases, obtaining suitable datasets poses significant challenges.

Unlike human medical datasets, which benefit from more standardised and widely available sources, veterinary datasets are relatively scarce. The images needed for training models in this domain must capture a variety of conditions across multiple breeds, often under nonideal circumstances. As a result, a custom dataset was developed to support this research, addressing the scarcity of available data. The following sections describe the DogEyeSeg4 dataset, constructed to overcome these challenges, as well as a synthetic dataset generated to augment the available real-world data.

3.1.1 DogEyeSeg4 Custom dataset

Given the limited availability of publicly accessible canine ophthalmic datasets, the development of the DogEyeSeg4 dataset became essential. The process of gathering suitable data from real-world clinical environments posed several difficulties. First, patient compliance during eye examinations often impacted image quality. Dogs, being non-cooperative subjects, frequently moved during assessments, resulting in blurry images or requiring multiple attempts to capture usable data. Second, the diversity of breeds presented additional variability in eye structure, fur colouration, and size, all of which affected the clarity and focus of the images. These factors, while reflective of real-world conditions, made it difficult to obtain consistent, high-quality images that were necessary for robust model training.

The images in the DogEyeSeg4 dataset were collected from two specialised veterinary ophthalmology clinics and a veterinary eye disease atlas [53]. Clinical environments, unlike controlled laboratory settings, introduce uncontrollable variables such as inconsistent lighting, non-standardised camera equipment, and varying angles of capture. These factors result in natural but challenging conditions for machine learning models. The lighting in veterinary clinics often varies depending on the location of the examination, leading to different levels of contrast and exposure in the images. Additionally, images were captured without staged conditions, meaning the dataset includes a range of natural scenarios rather than perfectly lit or artificially enhanced images. This variety adds to the complexity but also enhances the dataset's applicability to real-world diagnostic conditions.

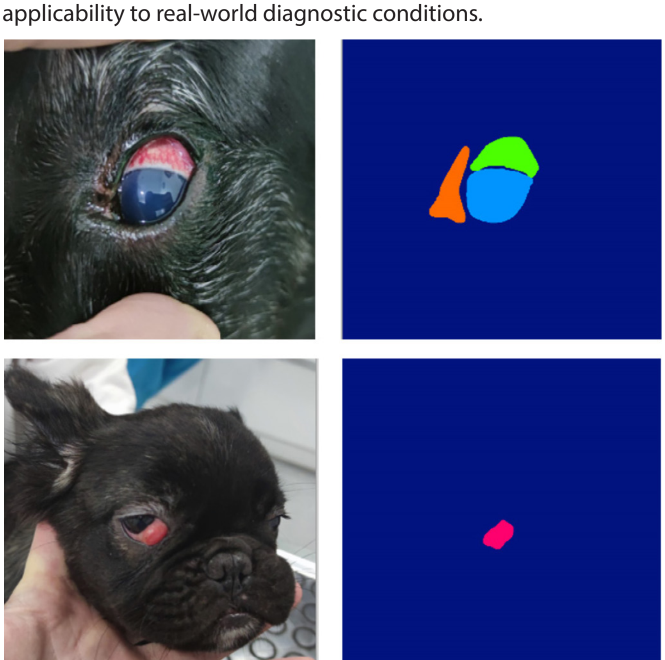

Figure 1. Example of images in DogEyeSeg4 dataset with corresponding masks showing closeup of an eye (upper row) and whole head (lower row)

The DogEyeSeg4 dataset consists of 145 images, which include both close-up images of the canine eye and full headshots. The example images with corresponding masks are shown in **Figure 1**. This diversity is crucial because veterinary practitioners often capture images that show either the entire head of the animal or focus specifically on the affected eye, depending on the diagnostic requirement. Close-up images provide detailed views of specific conditions, such as corneal cloudiness or excessive tearing, while whole-head images are more common in clinical settings and may capture multiple symptoms, including redness of the sclera or colored masses in the corner of the eye, from a broader perspective. Each image was resized to 320x320 pixels for consistency during model training, and the dataset is annotated with one-channel masks in PNG format.

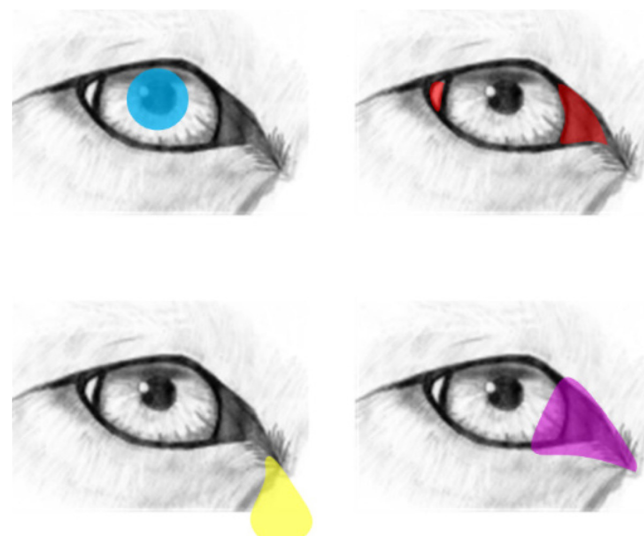

Figure 2. Visual representation of medical symptoms in the top-left to bottom-right order: S1 - cloudiness or haziness of the cornea; S2 - sclera redness; S3 - excessive tears; and S4 - colored mass protrusion in the corner of the eye. Precision of the dataset, especially given the variability in patient behaviour and the nonstandardized clinical conditions under which the images were captured

The four annotated classes, described in **Figure 2**.[54], correspond to the following symptoms:

- S1: Cloudiness or haziness of the cornea,
- S2: Redness of the sclera,
- S3: Excessive tearing,
- S4: A colored mass in the corner of the eye.

These symptoms correspond to several common diseases, such as Cherry Eye, Glaucoma, Uveitis, Corneal Ulceration, and Bacterial Keratitis. To ensure the clinical relevance and accuracy of the dataset, each image and its annotation were reviewed by a veterinary specialist. This review process was critical in maintaining the diagnostic.

Ensuring the dataset complied with data protection regulations, particularly the General Data Protection Regulation (GDPR) [55], was a critical part of the dataset creation process. All images included in the dataset were anonymised to safeguard the privacy of clients and their animals. Anonymisation involved the removal of any identifiable information, such as examination dates, client names, and animal identifiers. By adhering to GDPR guidelines, the dataset was rendered compliant with strict privacy regulations. Additionally, since the images were gathered during standard veterinary care and no animals were harmed for the purpose of image acquisition, ethical approval was not required for this study. This was made explicit in the ethics and consent statement accompanying the dataset: "Images were collected as part of routine clinical evaluations, and no ethical approval was necessary as no harm was caused to the animals."

One of the key challenges in assembling this dataset was balancing the representation of each class. For example, some symptoms, such as excessive tearing (S3), appeared more frequently than others, such as the colored mass (S4). Without careful selection, this imbalance could lead to biased model training, where the model becomes overly tuned to detecting the more frequent symptoms. Therefore, a meticulous curation process was employed to ensure a balanced distribution across all four classes.

To further enhance the dataset's utility, data augmentation techniques were applied. Building on methodologies from medical image segmentation, particularly with respect to U-Net architecture, augmentation methods such as horizontal flipping, rotation by up to 15 degrees, and translation by 50 pixels were used as seen in **Figure 3.** [56]. These augmentations introduce variation in camera angles and positions, simulating different real-world scenarios where veterinary professionals might capture images from slightly different perspectives. Importantly, zoom augmentation was avoided to prevent interpolation, which could introduce noise into the masks and degrade the accuracy of the annotations. The augmentation process expanded the dataset to 200 images, increasing its robustness and making it more suitable for training machine learning models that need to generalise across diverse imaging conditions.

Figure 3. Example of applied augmentations on original image in top row: horizontal flip (middle row left), horizontal shift (middle row right), rotation (bottom row left), and vertical shift (bottom row right)

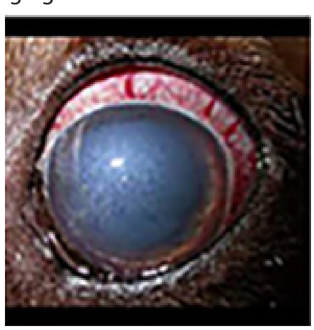

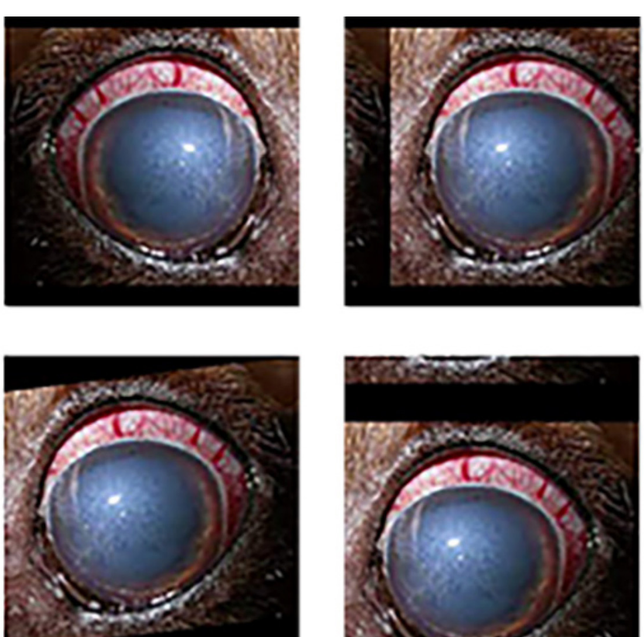

3.1.2 Synthetic datasets

Due to the limited availability of real-world data and the need for large datasets to train deep learning models, synthetic image generation has become a valuable resource in medical imaging, including veterinary ophthalmology. Several techniques exist for generating synthetic images, each offering strengths and weaknesses in tackling this issue. These methods include Generative Adversarial Networks (GANs) [57], Variational Autoencoders (VAEs) [58], and Diffusion Models [59].

Additionally, procedural methods like rule-based systems and 3D rendering techniques are also applied for synthetic image generation, offering high control over image features, albeit with increased manual intervention and resource demand [60].

GANs utilise a two-part system, where the generator creates images, and the discriminator attempts to distinguish them from real images. While GANs are capable of producing high-resolution, realistic images, they are computationally expensive and often unstable during training, with issues like "mode collapse," where limited variations of images are generated. VAEs, by contrast, are probabilistic models that encode input data into a latent space and then decode it back, generating new data by sampling from this latent space. VAEs are more stable to train than GANs but tend to produce lower-resolution images that may lack the detailed features required in tasks such as ophthalmic disease diagnosis. Diffusion models, particularly the Stable Diffusion variant used in this study, offer a more balanced approach, combining computational efficiency with high-quality image generation.

A significant feature of diffusion models is their reliance on the U-Net architecture. In diffusion models, the U-Net serves as the backbone for the denoising process, transforming random noise into coherent images. This structure makes diffusion models particularly well-suited for generating structured medical images [61].

Beyond learning-based techniques, rule-based image generation and 3D rendering techniques are also used in certain fields for procedural image creation [60]. Rule-based systems apply predefined algorithms to generate images, offering high customisation but requiring extensive manual intervention. Rendering techniques simulate 3D environments and lighting to create highly detailed images, but these methods demand considerable setup and computational resources [62].

Figure 4. Stable diffusion generated images with fixed posture of different dog breeds. Images shows high degree of reality with certain difficulty like double set of ears

For this study, Stable Diffusion was selected due to its balance between image quality and computational efficiency. Example images generated using Stable Diffusion can be observed in **Figure 4**. [63]. Diffusion models, particularly those built with the U-Net architecture, offer significant advantages in terms of structural detail and precision, which are crucial in medical applications like ophthalmology. The U-Net backbone, with its skip connections, helps preserve the fine-grained information necessary to capture subtle disease symptoms in canine eye images [18].

Stable Diffusion was further enhanced using Low-Rank Adaptation (LoRA), which allowed for parameter-efficient fine-tuning of the model, reducing both time and computational resources [64]. By using a small subset of real-world images from the DogEyeSeg4 dataset, Stable Diffusion was fine-tuned to specialize in generating images that accurately depict canine ophthalmic diseases like Glaucoma, Cherry Eye, and Uveitis. Example of such images using LoRA are presented in **Figure 5**. [63]. While rule-based systems and 3D rendering methods provide high control over specific image features, their high manual effort and resource demands made them less practical for large-scale image generation compared to Stable Diffusion [62].

Figure 5. Close-up Stable Diffusion images using custom LoRA describing various medical conditions based on the DogEyeSeg4 dataset

In diffusion models, the process of image generation involves gradually transforming random noise into coherent images. This denoising process is managed by a U-Net architecture, which effectively captures both high-level structures and fine details. The model is trained to learn how to reverse the noise-adding process, allowing it to generate clean, high-quality images from noisy input during inference [59], [61]. In this study, the fine-tuning of Stable Diffusion was performed using custom LoRA, which reduced memory requirements by updating only a subset of the model's parameters [65]. The use of a U-Net backbone allowed the model to capture the intricate details of various eye diseases, making it particularly effective for generating clinically relevant images of canine ophthalmic conditions.

One of the key features of Stable Diffusion in this study was its ability to add disease symptoms to otherwise healthy eyes using inpainting [66], [67]. Inpainting allows for localised changes to specific regions of an image while leaving the rest of the image unchanged. This feature was valuable for introducing symptoms like scleral redness or corneal cloudiness into healthy eye images, which can be examined in **Figure 6**. [63]. Example image. Using prompts such as "cloudy cornea with red sclera," the model could generate localised disease manifestations, allowing for the creation of synthetic images that accurately represented various stages of disease progression. This capability enhanced the diversity of the dataset, ensuring that a wide range of symptoms and severities were represented.

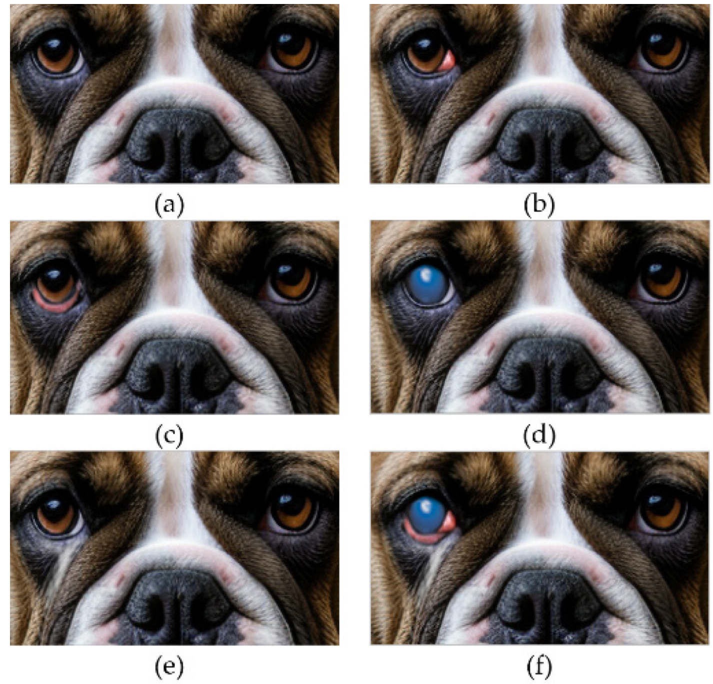

Figure 6. Stable Diffusion generated image containing symptoms using Inpainting showing (a) healthy eyes, left eye with: (b) prolapsed eyelid gland, (c) red sclera, (d) cloudy cornea, (e) epiphora and (f) all previously mentioned symptoms

In addition to generating new images, diffusion models offer unique opportunities for augmentation:

- **Symptom Severity Modification**: Adjusting the text prompts allowed for the generation of images with varying degrees of disease severity, from mild to severe symptoms [38].
- **Symptom Combination**: Diffusion models can create images with multiple symptoms, such as excessive tearing and a colored mass, replicating complex real-world cases.
- **Stochastic Variability**: The inherent randomness in diffusion models ensures that even with the same prompt, slight variations occur in the generated images, further diversifying the dataset without requiring additional real-world data collection [37].

Advantages of using Synthetic images:

- **Scalability of Data**: Synthetic image generation facilitates the creation of extensive datasets, which are particularly valuable for rare conditions or underrepresented canine breeds. This approach addresses the scarcity of real-world data, allowing researchers to simulate a wide range of clinical scenarios that might otherwise be difficult to capture [68].
- **Efficiency in Cost and Time**: Compared to the collection and annotation of real-world images, synthetic data can be generated rapidly and at a significantly lower cost. This enables the efficient scaling of datasets required for training deep learning models, reducing the time and resources associated with manual data acquisition [69].

>>>14

- **Controlled Variability**: Synthetic generation allows researchers to precisely control the inclusion of specific symptoms, conditions, and their severity. This level of control ensures that the dataset remains balanced, mitigating issues related to class imbalance, and comprehensively covering the spectrum of disease presentations [69].
- **Ethical Benefits**: The use of synthetic images circumvents the need for invasive clinical procedures or additional veterinary visits. As a result, it provides an ethically sound method for expanding datasets without subjecting animals to unnecessary tests or discomfort[40].

Limitations when synthetic images are used:

- **Realism Constraints**: Although synthetic image generation has made significant advancements, the resulting images may still exhibit subtle artefacts or unrealistic features. These imperfections could lead to a degradation in model performance when applied to real-world scenarios, as the models may struggle to generalise from synthetic to actual clinical data [39].
- **Bias Propagation**: Synthetic datasets, while generated artificially, can inadvertently carry over biases from the real-world data used in the finetuning process. This issue may limit the generalizability of models trained exclusively on synthetic data, as the diversity and complexity of real-world cases might not be fully captured [36].
- **Necessity for Clinical Validation**: Despite their utility, synthetic images require thorough validation to ensure their clinical relevance. Without rigorous validation processes, models trained on synthetic data might underperform when deployed in real-world clinical settings, especially when tasked with recognising nuanced or rare conditions [40].

3.2 Model Architecture and Training

3.2.1 U-Net

The U-Net architecture is widely regarded as a robust model for image segmentation, particularly in the field of biomedical image analysis. U-Net excels at pixel-wise classification tasks. Its architecture consists of two symmetric parts: an encoder (contracting path) and a decoder (expanding path), forming a U-shaped structure that allows the network to both extract features and reconstruct spatial details at the pixel level. The encoder in U-Net serves to downsample the input image through a series of convolutional and max-pooling layers, extracting increasingly abstract features. Each convolutional block consists of two 3x3 convolutional layers followed by a ReLU activation function. The downsampling occurs through max-pooling layers, which reduce the spatial resolution by a factor of two at each step. This progression allows the U-Net to capture the local features in early layers and more complex global patterns in deeper layers. Visual representation of the U-Net architecture is shown in **Figure 7**. [56].

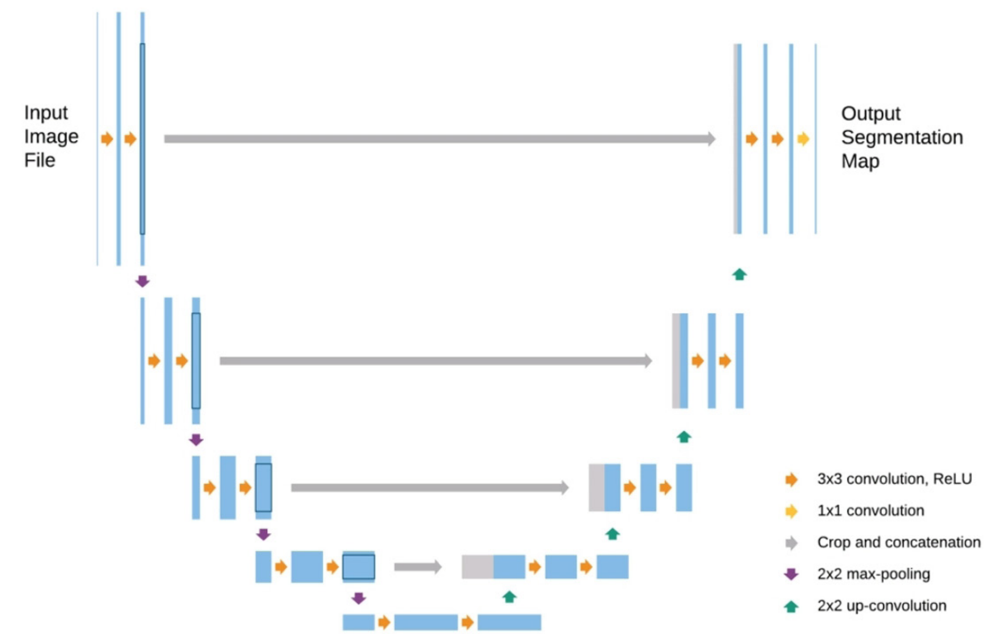

Figure 7. The standard U-Net architecture, where the left encoder part and right decoder part form the letter U. The encoder effectively captures detailed features from the input images, which are then upsampled by the decoder part of the U-Net. The decoder uses transposed convolutions to restore the image to its original size, using skip connections from the corresponding encoder layers to refine the segmentation output with high-resolution details retained [56]

At the deepest point in the architecture, the bridge links the encoder and decoder by combining features learned from the encoder's deepest layer with features from the decoder's shallowest layer. This design ensures that global and local features are propagated through the network, facilitating accurate segmentation even in cases of complex images with small or subtle regions of interest [70]. The combination of these feature maps is essential for generating detailed and contextually accurate segmentations, particularly in medical applications where high precision is critical. The decoder mirrors the encoder, progressively upsampling the feature maps and restoring the spatial resolution of the input image. The decoder combines the upsampled feature maps with the feature maps from the corresponding encoder layers via skip connections. These skip connections help retain the spatial context and fine details from the encoder, ensuring that the segmentation is accurate at a pixel level [18], [71].

The decoder uses transpose convolutions (deconvolutions) to increase the image resolution and further refine the output segmentation mask [72]. The final output is a pixel-wise segmentation mask, where each pixel is assigned a class label based on the features extracted by the encoder and decoder. U-Net has proven to be effective in tasks that require accurate segmentation with limited training data, making it particularly well-suited for medical imaging, where annotating large datasets is often challenging [18], [71]. Its symmetrical structure and use of skip connections make it highly capable of handling both large and small objects in an image, and it performs well even when there is significant variation in object size and shape.

3.2.1.1 Transfer learning

Transfer learning significantly enhances the capability of deep learning models, particularly in specialised tasks such as medical image segmentation, by leveraging pre-trained neural networks. The U-Net architecture benefits from incorporating robust feature-extracting networks, known as backbones, which are pre-trained on large and diverse datasets like ImageNet [73]. These backbones capture a wide array of features, from simple textures to complex structures, which are crucial for accurate segmentation.

One popular choice for a backbone in U-Net is the VGG network [31], known for its simplicity and effectiveness in feature extraction due to its deep architecture of 16 convolutional layers. The typical VGG-16 architecture, as shown in **Figure 8**. [56], consists of multiple convolutional layers followed by max-pooling layers, which progressively reduce the spatial dimensions while increasing the depth, capturing finer details at each layer.

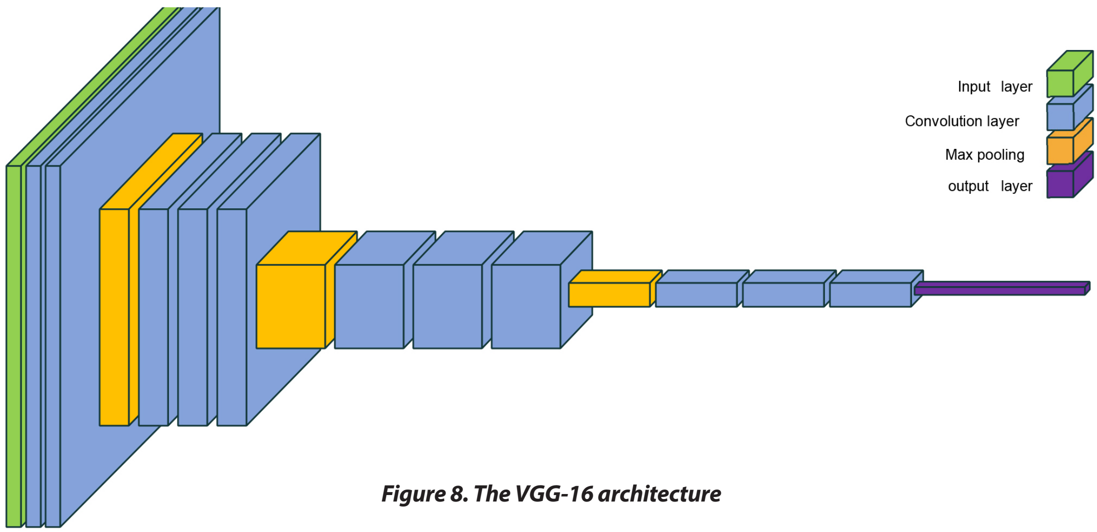

Figure 8. The VGG-16 architecture

Another powerful backbone is the ResNet18 [74], which incorporates residual connections to allow training of much deeper networks by addressing the vanishing gradient problem. These residual connections act as shortcuts that enable the gradient to flow through the network without diminishing, preserving the strength of the signal (**Figure 9**. [56]). This architecture is particularly beneficial for U-Net, enhancing its ability to learn from medical images where preservation of spatial hierarchies is critical.

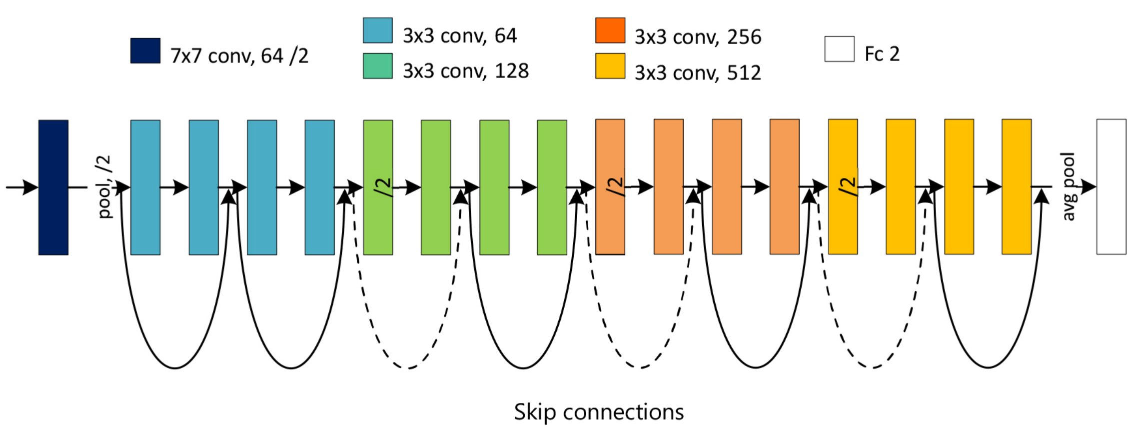

Figure 9. The ResNet18 architecture

The Inception V3 network [75], known for its efficiency in handling multi-scale information, uses modules composed of parallel convolutional layers with varying kernel sizes. This design allows the network to capture features at various scales within the same layer, making it highly effective for tasks that require the detection of objects of different sizes, such as different types of tissues or cells in medical images. The Inception V3 network is shown in **Figure 10**. [56].

>>>16

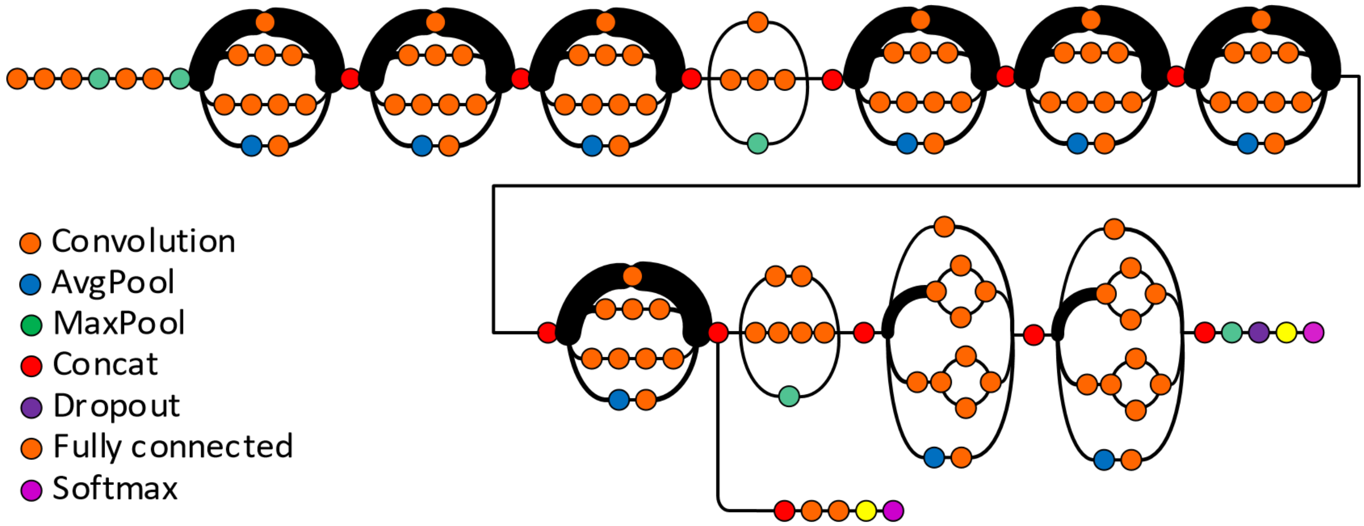

Figure 10. The Inception V3 architecture

Improved U-Net with Backbone Integration: The improved U-Net model incorporates these backbone networks into its architecture, where each encoder block is an instance of either VGG-16, ResNet18, or Inception V3. This integration allows the U-Net to utilise the advanced feature extraction capabilities of these networks, enhancing its performance significantly. The decoder of the improved U-Net model up-samples the feature maps and combines them with the extracted features at corresponding levels from the encoder, optimising the segmentation accuracy as seen in **Figure 11**. [56].

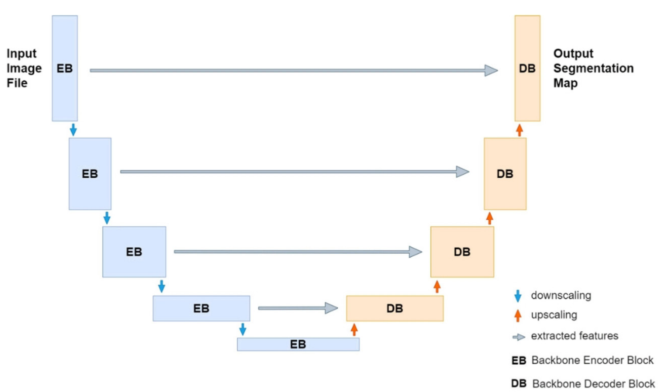

Figure 11. Schematic of the enhanced U-Net architecture with a backbone in both encoder and decoder blocks. The encoder blocks (EB) downscale the input while extracting features, and the decoder blocks (DB) upscale these features to reconstruct the output image

3.2.1.2 U-Net Algorithm

The training of the U-Net model involves an iterative process of adjusting network weights to minimise the loss between the predicted segmentation masks and the ground truth. The algorithm begins by initialising the weights of both the encoder and decoder networks randomly. The model is then trained using a dataset consisting of training images and corresponding ground truth segmentation masks.

For each training image, the following steps are performed in the training loop:

1. The encoder network encodes the input image, extracting a set of feature maps at different levels of abstraction.
2. The feature maps from the encoder are concatenated with the corresponding feature maps from the bridge, which combines contextual information from both high- and low-level layers of the network.
3. The concatenated feature maps are then passed through the decoder network, which progressively upsamples the feature maps and reconstructs the spatial resolution to generate a predicted segmentation mask.
4. The loss between the predicted segmentation mask and the ground truth segmentation mask is calculated using the selected loss function (such as Dice loss or cross-entropy loss).
5. The encoder and decoder network weights are updated using an optimisation algorithm (e.g., stochastic gradient descent or Adam) to minimise the loss.

The algorithm continues to iterate through the training images, updating the network weights after each image and its

U-Net Algorithm
Input: Training images and corresponding ground truth segmentation masks **Output: trained U-Net model**
Initialise encoder and decoder network weights randomly.
Training Loop
For each training image and ground truth segmentation mask:
Encode the image using the encoder network.
Concatenate encoder feature maps with bridge feature maps
Decode concatenated feature maps using a decoder network to generate a predicted segmentation mask.
Calculate the loss between the predicted and ground truth segmentation masks
Update encoder and decoder network weights using an optimisation algorithm.
Convergence Check
If loss converges or reaches a satisfactory level: Exit training loop **Else: Continue training loop**

Algorithm 1: U-Net Algorithm for Segmentation [56]

corresponding mask are processed. After each training loop, a check is performed to evaluate the convergence of the model. If the loss converges or reaches a satisfactory level, the training process is terminated. Otherwise, the loop continues until the model reaches optimal performance as described in **Algorithm 1**.

Once the model is trained, it can be used to predict segmentation masks for new images. The prediction involves passing the input image through the encoder and decoder, generating a probability map for each pixel in the image.
The class with the highest probability is assigned to each pixel, resulting in the final segmentation output. This process ensures that the U-Net model learns to accurately segment images based on the features extracted during the training phase, providing robust pixel-wise predictions for medical imaging tasks.

3.2.1.3 Loss function

In deep learning, loss functions are essential components that assess the disparity between a model's predicted output and the actual ground truth values. A higher loss indicates greater error in the model's predictions, while a lower loss suggests improved accuracy. The principal objective of loss functions is to guide the model toward minimising prediction errors by adjusting its trainable parameters, such as weights and biases.

There are numerous loss functions available, each designed to reward or penalise a model differently, depending on the nature of the task. The choice of the most appropriate loss function is critical for achieving optimal performance. In this research, two primary loss functions were used: Dice loss (DL) and Cross Entropy (CE) loss, both of which are highly relevant for segmentation tasks in medical imaging.

Dice Loss is directly derived from the Dice Similarity Coefficient (DSC), which measures the overlap between the predicted segmentation and the ground truth mask.

The formula for Dice loss, which ranges between 0 and 1, is expressed as (1):

1 $$DL = 1 - DSC$$

where *DL* tends toward 1 in cases of poor overlap and 0 when the model's predictions are perfect. Since Dice loss emphasises the similarity between the predicted and actual segmentation areas, it is particularly effective when working with small, imbalanced regions of interest. It calculates both **precision**—the accuracy of positive predictions—and **recall**, which measures the coverage of actual positive samples. Both precision and recall are integral to the effectiveness of the classifier and are expressed as follows (2):

2 $$precision = \frac{TP}{TP + FP} \qquad recall = \frac{TP}{TP + FN}$$

Using these definitions, the Dice Similarity Coefficient is equivalent to the **F1score**, which is a harmonic mean of precision and recall. This harmonic mean ensures that the model penalizes low values of either precision or recall, thereby encouraging a balance between the two. The F1-score is calculated as (3):

3 $$F1 = \frac{1}{\frac{1}{precision} + \frac{1}{recall}} = \frac{2 \cdot precision \cdot recall}{precision + recall}$$

In this research, Dice loss was used to transform this metric into a continuous, differentiable function that can be applied during the training of the model. The formula for Dice loss (4), when applied to predictions and ground truth, is:

4 $$DL(y, \hat{y}) = 1 - \frac{2\sum y \cdot \hat{y}}{\sum y + \hat{y}}$$

Here, y represents the ground truth values, and *y* accent represents the predictions made by the model. Dice loss performs well when applied to image-level predictions rather than individual pixels, as it accounts for the overlap between the predicted and true segmentation masks.

Another loss function used in this study is the Cross-Entropy (CE) loss, which is widely regarded for its simplicity and effectiveness in training classification models. Unlike Dice loss, which focuses on the overlap between areas, CE measures the divergence between two probability distributions: the predicted probability and the actual label. CE is particularly useful when working with multi-class classification problems, and it can be adapted for binary tasks. The Binary Cross Entropy (BCE) loss is used for binary classification tasks and is given by formula (5):

5 $$BCE(y, \hat{y}) = -(y \cdot \log(\hat{y}) + (1 - y) \cdot \log(1 - \hat{y}))$$

In cases where multiple classes are present, Categorical Cross-Entropy (CCE) is used. CCE extends BCE to multiple classes, and the formula for CCE (6) is:

6 $$CCE(y, \hat{y}) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{i,c} \log(\hat{y}_{i,c})$$

Where *FF* represents the total number of pixels, *DD* is the number of classes, and $yy_{ii,cc}$ refers to the ground truth label for class ρρ at pixel ρρ. Similarly, $yy_{ii,cc}$ represents the predicted probability for each class at each pixel. CCE focuses on minimising pixel-level errors but can struggle in imbalanced class distributions where larger objects dominate the scene, affecting the segmentation performance of smaller objects.

To address this imbalance, Focal Loss (FL) was introduced. Focal Loss builds upon Cross-Entropy by focusing more on hard-to-classify objects, particularly those incorrectly classified or with low confidence. This is achieved by down-weighting well-classified objects and focusing the learning process on challenging examples. Focal Loss is defined as (7):

7 $$FL(y, \hat{y}) = \alpha(1 - \hat{y})^{\gamma} \cdot BCE(y, \hat{y})$$

Here, αα controls the weighting of different classes, and *γγ* is the focusing parameter, which determines the degree to which the model focuses on difficult-to-classify examples. *γγ* = 0 Focal Loss behaves like Binary Cross-Entropy. For multi-class classification tasks, BCE is replaced with CCE, and the parameter αα becomes a vector of class weights, while *γγ* is expressed as a matrix of probabilities for each class.

3.2.1.4 Activation Functions

Activation functions are an essential component of neural networks, determining the output of each neuron and introducing non-linearity into the network. This nonlinearity enables the network to learn complex patterns and make accurate predictions. In U-Net and similar convolutional neural networks, common activation functions include ReLU, Sigmoid, and Softmax [76].

The ReLU (Rectified Linear Unit) activation function (8) is defined as:

$$ReLu(x) = \begin{Bmatrix} 1\ x > 0 \\ 0\ x \leq 0 \end{Bmatrix} \quad (8)$$

ReLU is widely used due to its simplicity and effectiveness, as it avoids the vanishing gradient problem that can occur with other activation functions like Sigmoid [77]. However, ReLU can lead to dead neurons, where neurons become inactive and stop learning. Variants like Leaky ReLU and Parametric ReLU have been introduced to address this issue [78].

The Sigmoid activation function (9), defined as:

$$Sigmoid(x) = \frac{1}{1+e^{-x}} \quad (9)$$

It is often used in the final layer of binary classification models. It maps the output to a value between 0 and 1, making it suitable for probability estimation. However, Sigmoid suffers from the vanishing gradient problem, which can slow down training [79].

For multi-class segmentation tasks, Softmax is typically used in the final layer. Softmax normalises the output probabilities for each class, ensuring they sum to 1. It is defined as (10):

$$SoftMax(x_i) = \frac{e^{x_i}}{\sum_{j=1}^{n} e^{x_j}} \quad (10)$$

Where *xxii* represents the input to the ρρ-th class. Softmax is particularly useful for multi-class classification, as it allows the network to assign probabilities to multiple classes and select the most likely class for each pixel [80].

3.2.2 GPT-4 and GPT-4o

Large Language Models (LLMs) such as ChatGPT, Mistral, Gemini, Claude, and Llama have demonstrated significant potential in medical diagnostics by aiding in symptom analysis and preliminary diagnosis. Each of these models brings unique strengths to processing medical data. For instance, ChatGPT excels at handling complex medical data and producing fluent, context-rich responses, while Mistral is specifically designed for efficient query resolution. Similarly, Gemini adapts well to dynamic conversational settings, and Claude offers a well-rounded approach to general data analysis. Despite these advantages, most of these models are challenged when tasked with complex, multi-symptom medical scenarios, as their architecture primarily focuses on textual understanding rather than deep multimodal analysis [69]. The introduction of GPT-4 marks a notable advancement over these earlier models. While the previously mentioned LLMs excel at interpreting symptom descriptions and providing initial diagnoses, GPT-4 surpasses them with its ability to incorporate context-driven responses and integrate detailed multimodal inputs [81]. More importantly, its evolved version, GPT-4o, extends the capabilities of the original model by incorporating object detection and segmentation tasks, which are critical in image-based diagnostics. This enhancement makes GPT-4o particularly valuable for tasks requiring a combination of natural language processing and detailed image analysis, as seen in veterinary ophthalmology diagnostics.

The GPT-4 model, recognised for its robust performance in language-related tasks, has been innovatively adapted for image analysis, deviating from its primary design as a text-processing engine. This unconventional application involves using images as input, where the model is expected to generate outputs that align with the visual content observed, rather than textual data. This method leverages the model's latent ability to handle multimodal tasks despite its foundational design as a language model [82]. While GPT-4 exhibits capabilities in image classification, it inherently lacks direct functionalities for object detection and segmentation—key components in detailed image analysis. Addressing these limitations, the enhanced version, GPT-4o, introduces mechanisms to perform these tasks by generating segmentation masks akin to those in ground truth data of test images, marking a significant enhancement over its predecessor [83]. The core architecture of GPT-4 is built upon the Transformer model, specifically tailored for language generation but applicable to a broader range of sequence processing tasks.

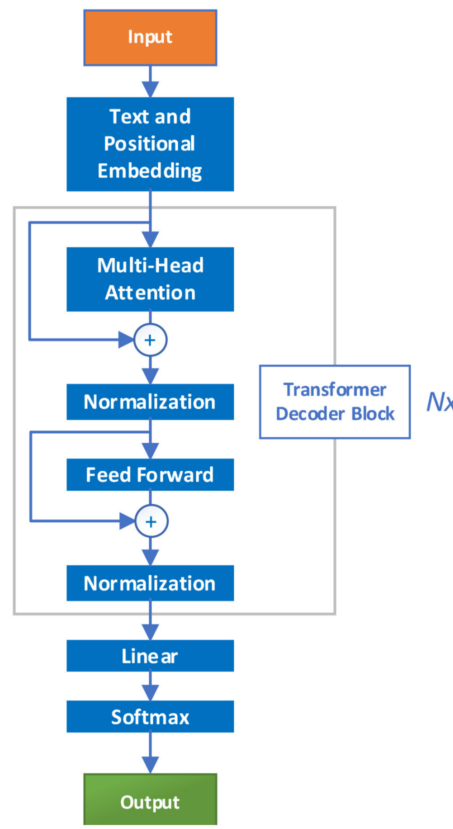

Figure 12. The architecture of a Transformer Decoder block is utilised in the GPT model. The process starts with embedding both the input tokens and their positional information. These embeddings are processed through several layers of multi-head attention and feed-forward networks, each layer followed by normalisation. The model's final output is produced by a linear layer followed by a softmax function. This design enables the model to attend to various parts of the input sequence concurrently, improving its ability to understand context and make accurate predictions

This architecture is illustrated in **Figure 12**. [63] and encompasses several critical components:

- Embedding Layer: Converts input tokens into rich, informative embeddings, supplemented with positional encodings to preserve the sequential nature of the input data.
- Multi-Head Attention: This mechanism allows the model to dynamically focus and refocus on different segments of the input data, facilitating a deeper understanding of the contextual interdependencies within the data.
- Feed-Forward Networks: Positioned to process the data sequentially from the attention mechanisms, these networks are pivotal in refining the data handling.
- Residual Connections and Layer Normalisation: These components are integral for maintaining the flow of gradients during training, thus enhancing the stability and performance of the model.
- Output Layers: Comprising a linear layer followed by a softmax activation, these layers predict the likelihood of the next token in the sequence, based on the processed embeddings.

To effectively deploy GPT-4 and GPT-4o in veterinary ophthalmology, specific contextual adaptations are necessary. These models are configured to act as theoretical consultants, analysing visual data from images and predicting possible veterinary conditions. The setup demands a tailored prompt that provides the models with a context encapsulating the task of diagnosing based on visual cues alone, without supplementary background information. An exemplary prompt used in this scenario is demonstrated in **Figure 13**. [63], where the models' output includes diagnostic suggestions and segmentation masks, crucial for practical applications [84].

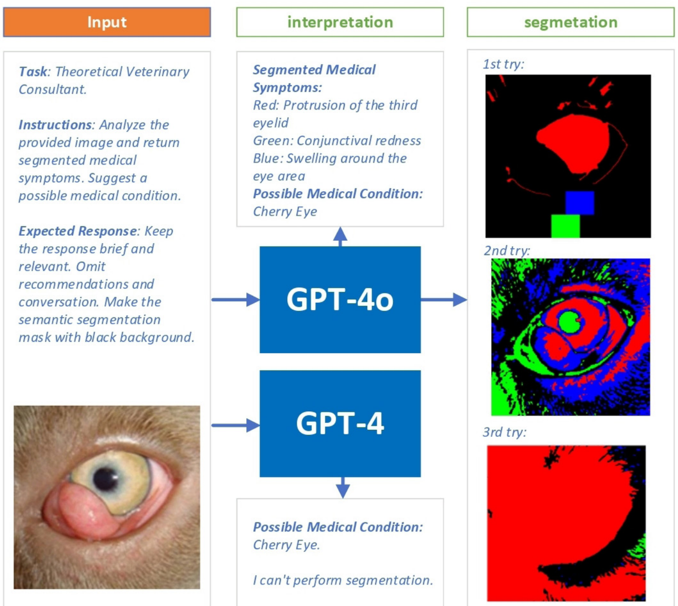

Figure 13. An example of an input prompt to GPT-4 and GPT-4o requesting segmentation of medical conditions and the most probable diagnosis

The outcomes derived from this innovative application of GPT-4 and GPT-4o are directly compared against traditional methods, providing a benchmark for assessing the efficacy of integrating language models in image-based diagnostic tasks. This comparative analysis is essential for validating the models' utility and effectiveness in real-world scenarios [85].

3.2.3 Grounding SAM

Object detection and segmentation models like YOLO [86] (You Only Look Once) and Mask R-CNN [87] (Region-Based Convolutional Neural Network) have shown remarkable performance in various applications, including medical image analysis. YOLO is known for its speed and efficiency, as it processes an image in a single pass to detect objects, making it highly suitable for real-time applications. However, both YOLO and Mask R-CNN require extensive training on customised datasets specific to the objects they are designed to detect or segment. This means that to achieve high accuracy in a particular domain, such as veterinary ophthalmology, these models must be fine-tuned using domain-specific data, which can be time-consuming and resource-intensive.

In contrast, Grounding SAM, a powerful combination of Grounding DINO [88] and SAM [89], operates in a more versatile and adaptable manner. Instead of relying on pre-training for specific objects, Grounding SAM only requires a text prompt to detect and segment objects in an image. This prompt-driven approach eliminates the need for pre-training on customised datasets, making it ideal for applications where labelled data is scarce or difficult to obtain. Grounding SAM builds upon the foundational principles of transformer models and leverages zero-shot learning capabilities, meaning it can generalise across different tasks and images without the need for extensive retraining. This characteristic sets it apart from traditional models like YOLO and Mask R-CNN, which depend on task-specific training.

Grounding DINO is the first step in the Grounding SAM framework, responsible for object detection. It operates by taking both images and text prompts as input. The transformer-based architecture of Grounding DINO enables it to process relationships between various objects in the image through its self-attention mechanism. This process allows the model to focus on relevant parts of the image, enhancing its ability to detect objects even in visually complex or noisy environments.

The text prompts provided to Grounding DINO are essential for guiding the object detection process. For instance, in the case of diagnosing canine eye diseases, text prompts describing symptoms in layman's terms (e.g., "extensive tearing") are used, as medical terminology might not yield effective results due to the model's limitations in domain-specific vocabulary.

Once the objects have been detected by Grounding DINO, the bounding boxes are passed to SAM for instance segmentation.

SAM's task is to convert these bounding boxes into precise segmentation masks, which provide pixel-level accuracy for each detected object. SAM utilises a transformer-based image encoder that can handle a wide range of segmentation tasks with minimal human intervention, making it highly versatile for medical applications, where segmentation accuracy is paramount.

In veterinary ophthalmology, SAM's segmentation masks are particularly useful for delineating fine details of ocular symptoms, such as corneal cloudiness or scleral redness. This level of precision in image segmentation ensures that even subtle abnormalities can be detected and analysed, providing critical data for diagnostic purposes. SAM's ability to generate accurate segmentation masks in a zero-shot manner—without the need for task-specific training—positions it as a powerful tool in the medical domain.

The combined operation of Grounding DINO and SAM is depicted in **Figure 14**. [63], which illustrates the process of diagnosing canine ocular conditions. The input image, typically a clinical photograph of a dog's eye, is first processed by Grounding DINO, which identifies areas of interest based on the provided text prompt.

Grounding DINO generates bounding boxes around the detected areas, such as the eye or tear ducts, and these are passed to SAM for detailed segmentation. SAM then produces segmentation masks that outline the detected regions with pixel-level precision. The segmented regions, such as areas of corneal opacity or excessive tearing, are used as input for further analysis.

A significant challenge in applying Grounding SAM to medical image analysis is prompt engineering, particularly when dealing with complex medical terminology. Grounding DINO and SAM requires prompts that can be interpreted accurately by the model.

Medical terms such as "epiphora" or "keratitis" may not yield effective segmentation results due to their specialised nature.

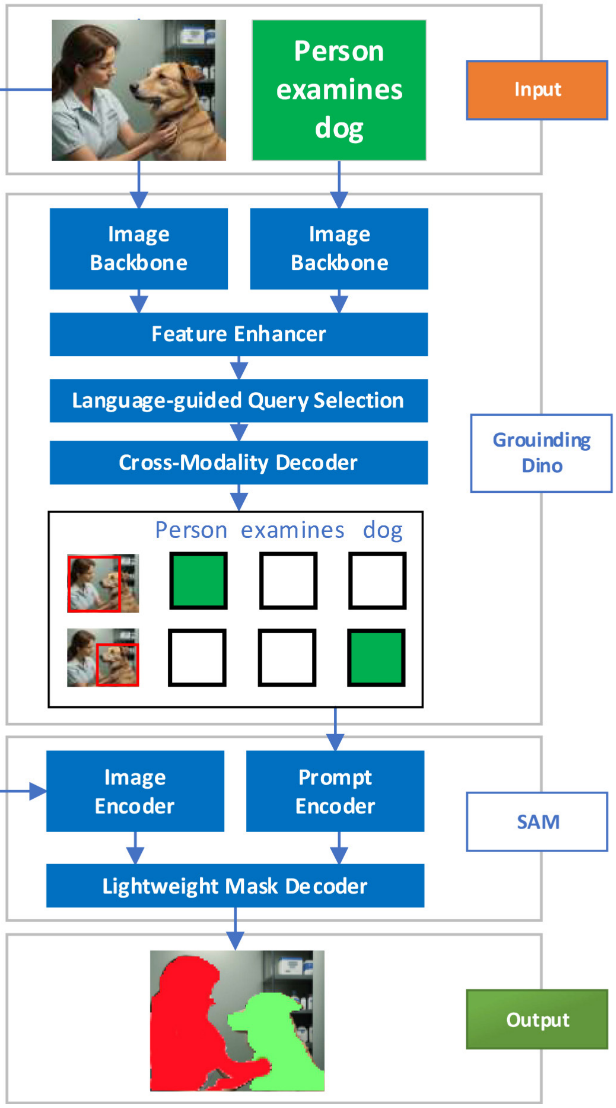

Figure 14. Grounding Dino and SAM integration flowchart

Instead, prompts need to be rephrased in simpler terms, such as "watery eyes" or "eye inflammation," to ensure the model can process and detect the relevant symptoms. **Figure 15**. [63] demonstrates how the structure and content of the text prompt significantly influence the model's output, highlighting the importance of careful prompt formulation.

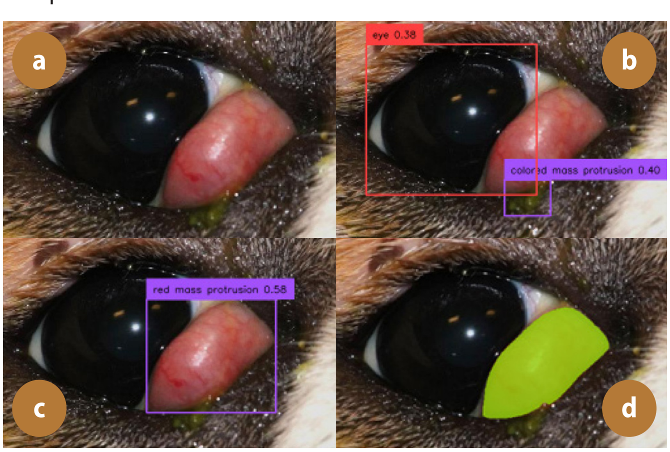

Figure 15. Grounding Dino prompting examples showing diversity with resulting detection and its prediction score. Prompts used for generating detection from top to bottom are: (a) "colored mass protrusion in the corner of the eye", (b) "colored mass protrusion in the eye" and (c) "red mass protrusion". Third prompt (c) is used for segmentation by SAM implementation (d).

3.2.4 Single Shot Multibox Detector (SSD) for Eye Detection

The Single Shot Multibox Detector (SSD) [90] was utilised to pinpoint the eye region in canine images, serving as an assisting tool for the improved segmentation task. SSD stands out for its real-time performance, accomplishing object localisation and classification simultaneously in a single forward pass of the network. This efficiency is achieved by omitting a separate region proposal generation stage, a common feature in other detection systems like Mask R-CNN, which simplifies the detection process and reduces computational overhead.

The architecture of SSD is built upon a robust base convolutional neural network, which is employed for feature extraction from the input image. Following this feature extraction phase, SSD utilises additional convolutional layers to produce feature maps at multiple scales. This multi-scale feature extraction allows SSD to detect objects of various sizes effectively. Key components of the SSD workflow include:

- **Feature Extraction**: which utilises a pretrained network like VGG16 or ResNet to derive rich feature maps from the input image. These features form the foundation for detecting objects at various scales.
- **Detection Heads**: a segment where each feature map level has associated detection heads that output scores and bounding boxes for potential objects at that scale.
- **Non-Maximum Suppression (NMS)**: this step refines the detection by removing overlapping boxes, ensuring that each detected object is represented by the single best bounding box.

The process flow of the SSD model is illustrated in the provided diagram in **Figure 16**, which clarifies the stages from image input through to the final predictions. The diagram shows how the image is processed through the VGG16/ResNet layers for feature extraction, followed by multi-scale detection heads, and culminating in the NMS stage to produce the final object predictions.

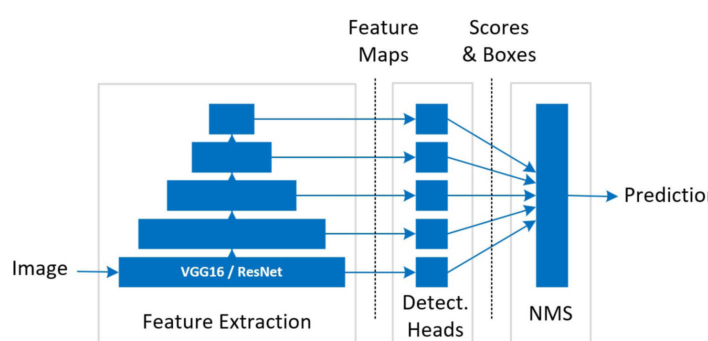

Figure 16. Single Shot Multibox Detector diagram for eye detection

Given the absence of publicly available models for detecting canine eyes, a custom SSD model was trained specifically for this task. The training dataset consisted of 99 images of dog eyes from various breeds and medical conditions, including healthy eyes, each annotated with ground truth bounding boxes in XML format. This approach was designed to ensure that the model could accurately identify eye regions across diverse scenarios.

The performance of the custom-trained SSD model was evaluated with impressive outcomes. The model achieved an average Intersection over Union (IoU) of 0.92 with a standard deviation of 0.08, and an average Dice Score of 0.96 with a standard deviation of 0.05. An F1 score of 0.9, calculated using Precision (0.97) and Recall (0.84), defines a well-suited detection model. These results highlight the SSD model's high precision in localising the eye region within the canine images, confirming its effectiveness for this specific application in veterinary ophthalmology. While the model demonstrates strong performance metrics, occasional detection inaccuracies underscore the need for ongoing adjustments and improvements in model training.

3.3. **Web Application Development**

The development of the web application for eye symptoms segmentation and medical condition prediction focuses on providing a robust, user-friendly interface to assist veterinarians and pet owners in diagnosing canine ocular diseases using AI-based image analysis. The web application is built to integrate a suitable segmentation model, ensuring seamless deployment, accessibility, and scalability. The development process is initiated by designing an intuitive user interface that would guide users through the image acquisition and diagnostic process. As seen in **Figure 17**, the system starts with a user interface that allows users to either capture images of their dog's eyes or upload them from local storage.

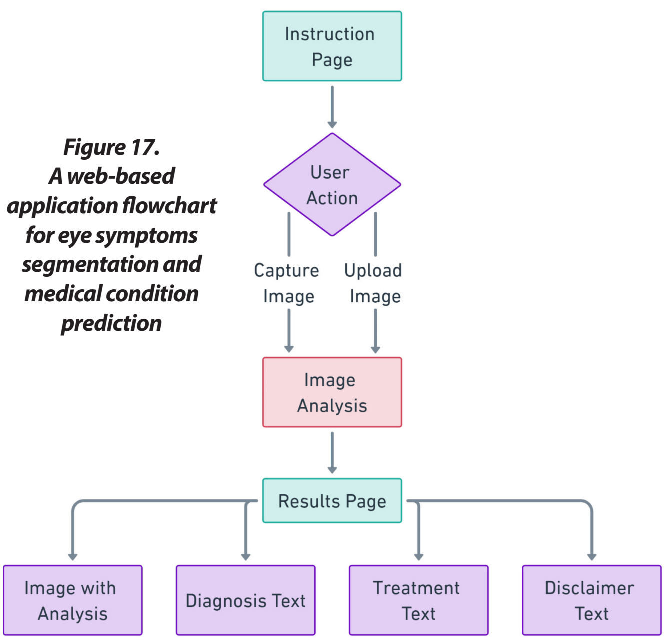

Figure 17. A web-based application flowchart for eye symptoms segmentation and medical condition prediction

This flexibility accommodates both veterinary professionals and pet owners who may not have prior experience with such diagnostic tools. The quality and framing of these images are crucial for accurate segmentation [56], so the interface should include an overlay mask as shown in **Figure 18**.[54], guiding users on how to correctly align the dog's eye for optimal image capture [54].

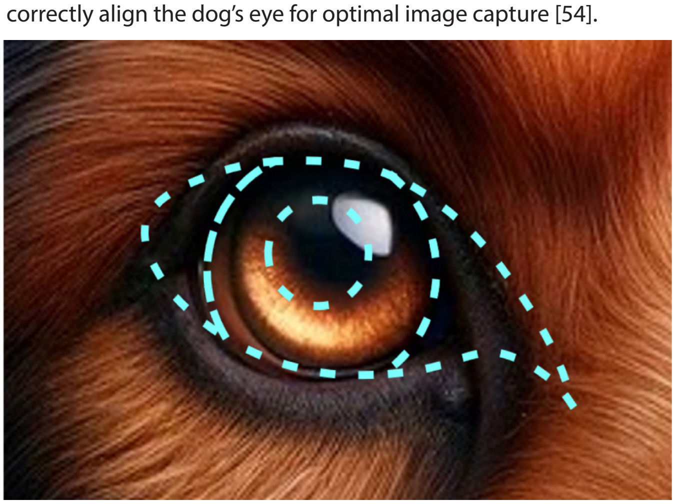

Figure 18. An overlay mask acting as a capturing images guide for accurate segmentation

The backend architecture of the application was designed to support the high computational demands of the segmentation model while ensuring portability. The system is containerised using Docker [91], which simplifies the deployment process and ensures that the application runs consistently across different platforms. This containerised approach also allows the application to scale efficiently, accommodating an increasing number of users without sacrificing performance. The core image processing tasks are handled by a suite of Python [92] libraries, including TensorFlow [93] for model execution, OpenCV [94] for image manipulation, and Pillow [95] for image formatting. The backend server, built using Flask [96], ensures stable communication between the web interface and the desired segmentation model, processing images and delivering results in real-time [54]. The entire setup ensures that the application is portable, scalable, and easy to deploy on both local machines and cloud-based servers.

To ensure efficient handling of requests and model processing, the deployment was optimised for server environments equipped with advanced graphics processing units (GPUs), ensuring that the heavy computational tasks, such as segmenting and analysing canine eye images, are handled with minimal latency. This server-based approach also allows for future scalability, where additional images can be collected and processed, contributing to the improvement of the model through continuous retraining and dataset expansion. Additionally, the application is built with the capacity to expand into mobile platforms, further enhancing its accessibility.

The web application not only provides diagnostic insights but also aims to enrich the model by collecting user data with their consent. Each user interaction and image submission can potentially contribute to the growth of the dataset, allowing for further model refinement. This continuous loop of data collection, analysis, and improvement represents the application's potential for evolving alongside the veterinary community's needs. The user experience is further enhanced with feedback features, guiding users with next steps based on the results, and providing valuable insights into their pet's ocular health.

The combination of these technical and practical design choices has led to a comprehensive tool for diagnosing canine eye diseases, setting the foundation for further developments and improvements.

4. EXPERIMENTS AND RESULTS

This section presents a detailed analysis of the experimental procedures and outcomes, focusing on the performance of various models in diagnosing canine ocular diseases. It is structured to provide insights into the evaluation process through several stages. First, we describe the test dataset used for model evaluation, including its structure and content. Next, we elaborate on the metrics applied to assess both segmentation and text generation tasks.

The core of the experimentation focuses on the training and evaluation of the U-Net model, both with and without backbone networks. A comparative study is also performed between the U-Net model trained on a real dataset versus one trained on a combined real and synthetic dataset to evaluate the impact of synthetic data augmentation. In a separate experiment, the benefits of binary vs. multiclass models were assessed along with their limitations.

Additionally, this section compares the performance of U-Net against GPT-4 and Grounding Dino in both segmentation and diagnostic interpretation tasks.

>>>22

These comparisons highlight the strengths and weaknesses of different approaches, allowing for a comprehensive assessment of model efficacy in real-world applications.

4.1 Test dataset

The test dataset for evaluating the U-Net, GPT-4, and Grounding SAM models was meticulously assembled to ensure diversity and clinical relevance, drawing from a wide range of veterinary resources. The images in the dataset were sourced from veterinary medical encyclopedias, peer-reviewed medical articles, and specialised veterinary clinics, all of which provided documented diagnoses and necessary permissions [53]. Each image in the dataset was resized to a standardised 320x320 pixel format to facilitate consistent processing and analysis across models. This standardisation is crucial to maintain the uniformity of input data and ensure the comparability of model outputs during testing. A representative example from the dataset is shown in **Figure 19**. includes images of an American bulldog [97] and a crossbred dog [98], both diagnosed with Cherry Eye as seen in. These images depict a pink-colored ocular protrusion from the medial canthus—a classic sign of Cherry Eye—accompanied by excessive tearing. Although both dogs present similar symptoms, their differing breeds introduce subtle anatomical variations that challenge automatic segmentation models. The models must learn to handle such variations and correctly interpret symptoms in diverse breeds. This illustrates some of the inherent difficulties that computer vision must address when attempting to accurately segment and diagnose eye conditions in canines.

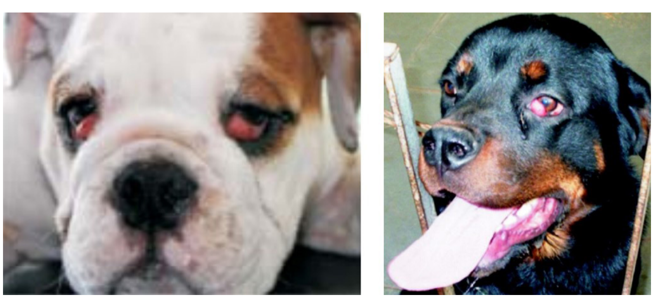

Figure 19. Images of an American bulldog (left) and a crossbred dog (right) show a pink-colored mass protruding from the medial canthus of both eyes, accompanied by signs of inflammation, including pain and excessive tearing, indicative of Cherry Eye

The test dataset was annotated in detail to reflect four key ocular symptoms that are commonly observed in veterinary ophthalmology: Corneal Cloudiness or Haziness, Sclera Redness, Excessive Tearing and Colored Mass Protrusion in the Corner of the Eye. Each image was paired with a one-channel mask, stored in PNG format, annotating the four distinct classes of symptoms. These annotations are crucial for training and validating the models, as they provide ground truth data that allows the models to learn the correct segmentation patterns. The pixel intensity in the mask varies depending on the symptom being represented, allowing for precise identification of each condition. This approach mirrors the methodology used for the DogEyeSeg4 custom dataset, ensuring consistency between training and test datasets.

The dataset includes images representing five significant ocular diseases in canines: Cherry Eye, Glaucoma, Uveitis, Corneal Ulceration, and Bacterial Keratitis. Each of these diseases is featured through five distinct cases, providing a total of 25 images for comprehensive model testing.

- **Cherry Eye**: This condition is characterised by the prolapse of the gland of the third eyelid, visible as a pink-colored mass in the corner of the eye. Cherry Eye can cause discomfort and, if left untreated, may lead to more severe complications such as dry eye or infection. The test images show various stages of Cherry Eye, providing the models with different levels of protrusion and tear production to segment and interpret.
- **Glaucoma**: Caused by elevated intraocular pressure, Glaucoma often results in pain, corneal cloudiness, and scleral redness. The dataset includes images showing a range of Glaucoma severity, from mild to advanced cases, where the eyeball is visibly enlarged. These images challenge the models to recognise subtle differences in redness intensity and cloudiness levels.
- **Uveitis**: This inflammatory condition affects the uveal tract and often presents with scleral redness, pain, and sensitivity to light. In some cases, corneal cloudiness is also observed. The Uveitis cases in the dataset emphasise the need for precise segmentation of redness and cloudiness, ensuring that the models can distinguish between inflammatory conditions and other ocular issues.
- **Corneal Ulceration**: This condition is typically the result of trauma or infection and is marked by corneal erosion or scratches. The test images highlight the symptoms of excessive tearing and sclera redness, along with visible cloudiness in some cases. These images require the models to focus on small, localised areas of damage, a challenge in real-world clinical scenarios where the ulcers may be difficult to detect without specialised equipment.
- **Bacterial Keratitis**: Bacterial infection of the cornea leads to symptoms such as corneal cloudiness and tearing and can progress rapidly if untreated. The dataset provides examples of varying severity, allowing.

The models to learn how to segment the infected areas accurately and distinguish them from less severe conditions like Corneal Ulceration.

Each disease class is represented with diverse images, ensuring that the models are exposed to a broad spectrum of cases. This enables the models to generalise better and handle real-world variations in symptom presentation, breed anatomy, and disease progression. These images, along with their corresponding annotated masks, form the foundation for the final model evaluation. The sample of the test dataset can be observed in **Figure 20**. [63].

This test dataset is not only a tool for evaluating the U-Net segmentation model but also serves as a benchmark for comparing the performance of the GPT-4 and Grounding SAM models.

By utilising the same dataset across all models, the results can be compared directly, allowing for a thorough performance evaluation. The models will be assessed based on their ability to accurately segment and interpret the symptoms associated with each disease, with particular focus on their precision in identifying small, nuanced differences between similar conditions.

The goal of this evaluation is to determine which model offers the most effective solution for automated segmentation and diagnosis in canine ophthalmology. The evaluation will be detailed further in the "Performance Evaluation" section, where we will outline the specific metrics and outcomes that support this comparative analysis.

« BACK TO CONTENTS

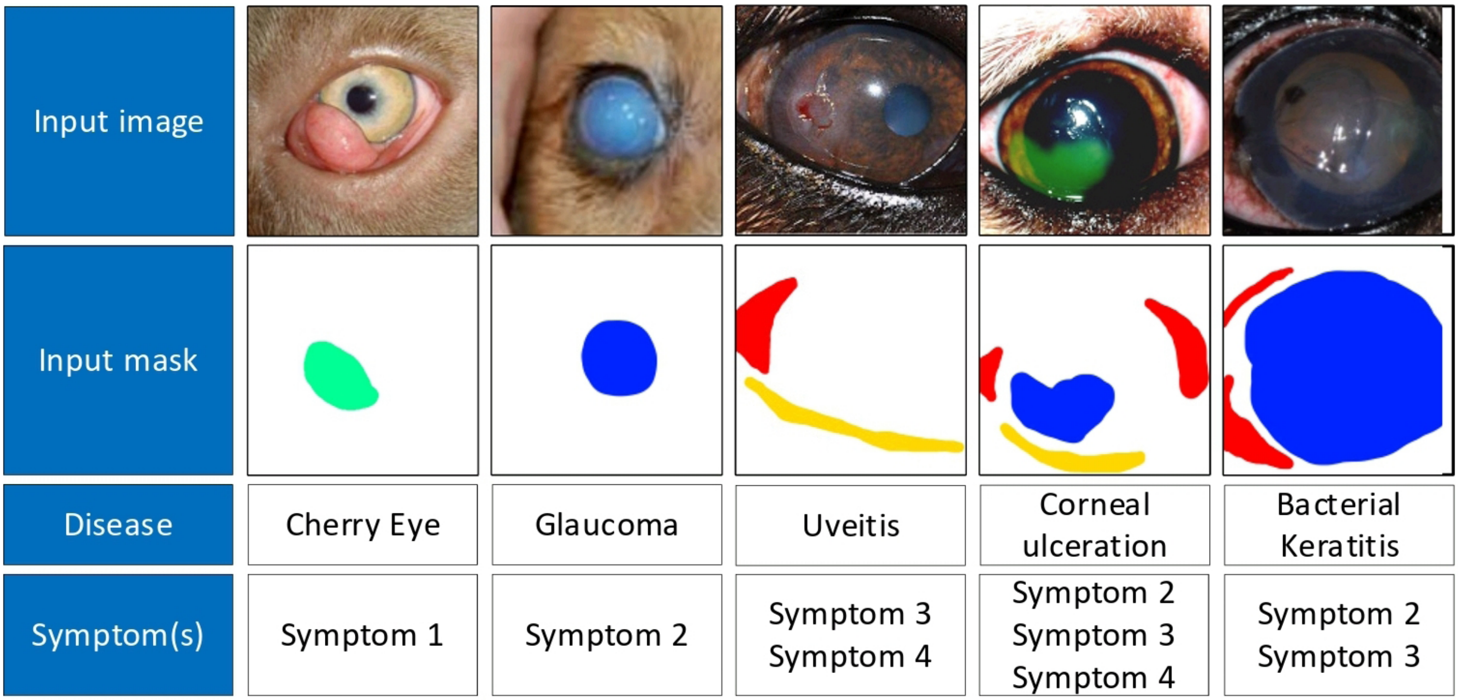

Figure 20. The test dataset contains 25 samples, evenly distributed across five different diseases. Symptom 1 represents a colored mass protruding in the corner of the eye, Symptom 2 refers to cloudiness or haziness in the cornea, Symptom 3 signifies redness of the sclera, and Symptom 4 denotes excessive tearing

4.2 Metrics

The evaluation of the experiment is centered around two primary aspects: the **segmentation of symptoms** and the **interpretation of input images**. The results from both stages of the experiment will be quantitatively assessed using standard metrics for segmentation accuracy and text similarity, comparing the model outputs to ground truth data. The following subsections provide detailed explanations of the metrics and procedures employed for evaluating model performance.

4.2.1 Segmentation metrics

To assess the effectiveness of the segmentation models, several well-established metrics are applied. These metrics offer a robust evaluation of how accurately the models segment the key symptoms in canine ocular images.

- **Jaccard Index** (Intersection-over-Union, IoU): The Jaccard Index measures the overlap between the predicted segmentation and the ground truth, reflecting how well the model can identify the same regions as the ground truth segmentation. It is calculated as (11):

$$\mathrm{IoU} = \frac{TP}{TP + FP + FN} \tag{11}$$

Where TP represents true positives (correctly identified regions), FP represents false positives (incorrectly identified regions), and FN represents false negatives (missed regions) as described in **Figure 21**. IoU provides a direct measure of how much of the predicted segmentation overlaps with the actual, intended segmentation.

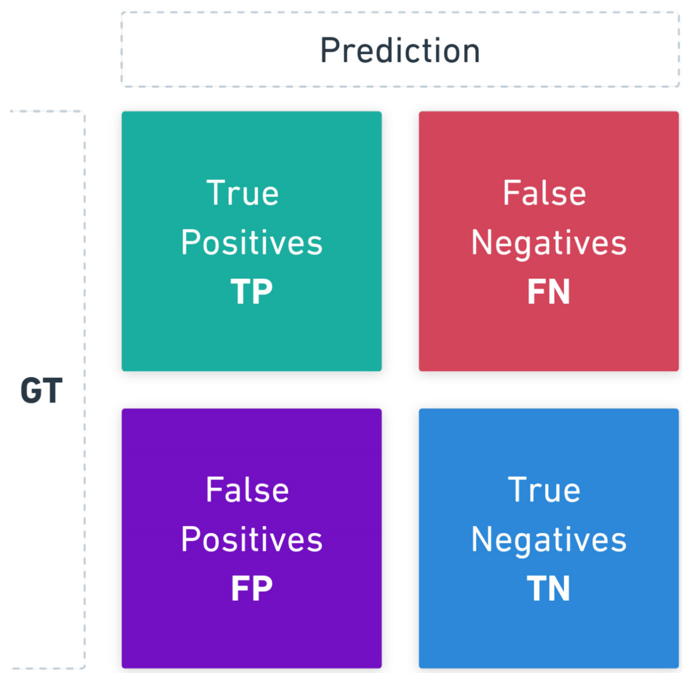

Figure 21. Representation of a confusion matrix with horizontal argument as Ground Truth (GT) and vertical as Prediction, showing correctly identified regions as TP, incorrectly identified regions as FP, missed regions as FN and correctly identified negative regions as TN

- **Dice Index**: A measure similar to the Jaccard Index but with greater sensitivity to overlap, the Dice Index evaluates the degree of similarity between the predicted and actual segmentations. It is calculated using the following formula (12):

>>>24

12
$$\text{Dice Index} = \frac{2 \times TP}{2 \times TP + FP + FN}$$

The Dice Index is particularly useful in medical imaging as it penalises small discrepancies between the predicted and actual segmentation more strongly than IoU, thus providing a more stringent assessment of segmentation performance.

- **Pixel Accuracy**: This metric calculates the proportion of correctly classified pixels across the entire image, both for foreground and background classes. The formula (13) is:

13
$$\text{Pixel Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Where TN refers to true negatives (**Figure 21**). While Pixel Accuracy offers a broad overview of the model's ability to classify pixels, it may be less sensitive to rare conditions when most of the image consists of healthy tissue.

- **Sensitivity** (Recall): Sensitivity, also known as the true positive rate, measures the proportion of actual positives (symptomatic regions) that the model correctly identifies. It is a critical metric in medical diagnostics because it indicates the model's ability to detect all possible symptomatic areas. The formula (14) is:

14
$$\text{Sensitivity} = \frac{TP}{TP + FN}$$

A high sensitivity score is crucial for ensuring that no symptomatic regions are missed.

- **Specificity**: Specificity, or the true negative rate, evaluates how well the model identifies negative regions, which are areas of the image that are free from symptoms. The formula (15) is:

15
$$\text{Specificity} = \frac{TN}{TN + FP}$$

Specificity is important for minimising false positives, ensuring that the model does not incorrectly classify healthy regions as symptomatic. Both Sensitivity and Specificity work together to provide a comprehensive evaluation, ensuring that the model not only identifies symptomatic regions but also avoids misclassifying healthy areas.

For all these metrics, a score closer to 1 indicates better performance, signifying near-perfect alignment between the predicted segmentation and the ground truth mask. These metrics will be used to compare the performance of U-Net with other models, providing a detailed analysis of how well the models generalise across different cases.

4.2.2 Text similarity metrics

In the second part of the experiment, the models' ability to interpret images and generate accurate textual diagnoses is evaluated. This is particularly important for models like GPT-4 and GPT-4o, which generate text-based interpretations of the symptoms seen in the images. A suite of metrics commonly used in natural language processing (NLP) is employed to compare the generated text with the reference diagnosis.

- **BERTScore**: BERTScore is a sophisticated metric that uses BERT embeddings to evaluate the semantic similarity between the generated text and the ground truth. It measures the cosine similarity between the tokens of the predicted and reference sentences (16):

16
$$BERTScore(P_B, R_B, F_B) = P, R, F1$$

where TT_{BB} and RR_{BB} represent the precision and recall of the BERT embeddings, respectively. BERTScore captures subtle semantic similarities that traditional n-gram-based metrics may overlook. The formulas for Precision (17) and Recall (18) are:

17
$$P_B = \frac{1}{|C|} \sum_{t \in C} \max_{t' \in R} \cos(emb(t), emb(t'))$$

18
$$R_B = \frac{1}{|R|} \sum_{t' \in R} \max_{t \in C} \cos(emb(t'), emb(t))$$

Where C represents the set of tokens in the candidate sentence and R the set of tokens in the reference sentence, the term *cosρρSSee (SS), ρρSSee (SS)* denotes the cosine similarity between the embeddings of tokens t and t'. The F1 Score (FF_{BB}) is computed as follows (19):

19
$$F_B = 2\frac{P_B R_B}{P_B + R_B}$$

- **CLIPScore**: Leveraging the multimodal CLIP model, CLIPScore measures the similarity between the vector representations of the generated text and the reference text. It is calculated using the cosine similarity between the two vectors (20):

20
$$CLIPScore = \frac{V_{candidate} V_{reference}}{\|V_{candidate}\| \|V_{reference}\|}$$

Where $VV_{ccccnncciiccccttcc}$ and $VV_{ccrrccrrccnncccc}$ represent the vector embeddings of the candidate and reference texts.

This metric is particularly effective for evaluating descriptions in the context of images, making it ideal for diagnosing medical conditions from image input.

- **BLEU**: This metric evaluates the overlap of n-grams (sequences of words) between the candidate and reference texts, penalising shorter sentences to ensure that full descriptions are considered. The formula is (21):

21
$$\text{BLEU} = \text{BP} \cdot \exp\left(\sum_{n=1}^{N} w_n \log p_n\right)$$

BLEU is a popular metric for evaluating machine translation and text generation models.

- **METEOR**: This metric computes the harmonic mean of unigram precision and recall, leveraging linguistic resources such as stemmers and synonyms. The formula is (22):

22
$$\text{METEOR} = \text{F}_{METEOR} = 10\frac{PR}{9P + R}$$

METEOR provides a more nuanced evaluation by considering word stems and paraphrasing.

« BACK TO CONTENTS

- **ROUGE**: ROUGE measures the overlap between the longest common subsequences in the candidate and reference texts, without requiring consecutive matching words. ROUGE-L is a common variant used (23):

23 $$\mathrm{ROUGE_L} = \frac{(1+ß^2)RP}{R+ß^2P}$$

where R is recall, P is precision, and β is a weighting factor.

- **SPICE**: SPICE compares the scene graphs of the candidate and reference sentences, focusing on the semantic content. It evaluates the propositional meaning of the sentences rather than just surface-level ngram matches (24):

24 $$\mathrm{SPICE} = \mathrm{F}_{SPICE} = 2\frac{PR}{P+R}$$

By employing these diverse and complementary text similarity metrics, we ensure a comprehensive evaluation of the models' ability to generate accurate, semantically meaningful diagnoses from medical images.

Each metric provides a unique perspective on the model's performance, capturing nuances in both lexical accuracy and semantic understanding. This multi-metric approach ensures that the models are evaluated rigorously, supporting the reliability of their medical interpretations.

5. DISCUSSION

This study was guided by two primary hypotheses:

Hypothesis 1: **A computer vision model can recognize certain canine ocular conditions in still images taken in an unconstrained environment**.

Hypothesis 2: **Modification of the input and architecture of the U-Net network contributes to a better segmentation of canine eye conditions**.

In the "Experiments and results" section, strong evidence supporting both hypotheses is provided. We developed and evaluated a series of models for detection, segmentation, and interpretation tasks to diagnose canine ocular diseases from images captured in real-world, unconstrained environments. Given the lack of pre-existing models for detecting canine eyes, we trained a custom Single Shot Multibox Detector (SSD) model. This SSD model was trained on 99 images encompassing various breeds and medical conditions, including healthy eyes. The model achieved the next scores: IoU 92%, Dice Score 96%, Precision 97%, Recall 84% and F1 Score of 90%.

These metrics indicate high accuracy in localising the eye region within unconstrained images, confirming that a computer vision model can effectively detect canine eyes in real-world conditions, which is needed for the experiment with a custom attention mechanism used by the U-Net segmentation.

For symptom segmentation, we initially implemented the standard U-Net architecture and then enhanced it with four different backbone networks: ResNet34, Inception V3, VGG16 and EfficientNet B3. These models were trained using the next loss functions: Categorical Cross-Entropy (CCE), Dice Loss (DL), Focal Loss (FL) and Combined Dice and Focal Loss (DL+FL). Best Performing Model is the U-Net model with a ResNet34 backbone trained using the combined Dice and Focal Loss (U-Net + ResNet34 with DL+FL). In 5-fold cross-validation, it achieved: Mean IoU 74%, Dice.

Similarity Coefficient (DSC) 83.4%, Sensitivity 86.3% and Pixel Accuracy 97.2%. The U-Net + ResNet34 with DL+FL outperformed other configurations in three out of four classes and demonstrated robust performance in accurately segmenting key ocular symptoms: S1 - Cloudiness or haziness of the cornea, S2 - Redness of the sclera, S3 - Excessive tearing and S4 - A colored mass in the corner of the eye.

Further enhancements to the segmentation model were made by training on a combined dataset of real images from DogEyeSeg4 and 400 synthetic images generated using Stable Diffusion, referred to as U-Net(RSD). This model achieved IoU 81%, DSC 88% and Pixel Accuracy 97%. These improvements underscore the effectiveness of augmenting training data with synthetic images to enhance model generalisation and performance.

The U-Net(RSD) model was compared with other segmentation approaches: GPT-4o (IoU 16%, DSC 20% and Pixel Accuracy 72%), Grounding SAM (IoU 14%, DSC 16% and Pixel Accuracy 67%). The U-Net(RSD) model outperformed these models across all metrics, highlighting its superior capability in segmenting canine ocular symptoms in unconstrained environments.

We explored training individual binary segmentation models for each symptom versus a single multiclass model. Binary models were trained using heatmaps as 4th channel in the RGB input image generated using information from the SSD eye detection.

While binary models with heatmaps reduced false positives compared to binary models trained without a custom attention mechanism, they faced challenges with overlapping symptoms and increased complexity. The multiclass U-Net(RSD) model provided better overall performance and operational efficiency, supporting its use for practical applications.

For diagnostic interpretation, several Large Language Models (LLMs) were evaluated: GPT-4, Gemini, Mistral, Claude and Llama-3.

GPT-4 consistently provided the most accurate diagnoses based on the segmented symptoms. For single-diagnosis scenarios, GPT-4 achieved: BERTScore 0.80, CLIPScore 0.91, BLEU 0.52, METEOR. 0.42, ROUGE 0.56 and SPICE 0.56.

When integrating segmentation masks from U-Net(RSD), the succeeding model GPT-4o demonstrated enhanced performance: BERTScore 0.84, CLIPScore 0.91 and BLEU 0.49, METEOR 0.32, ROUGE 0.51 and SPICE 0.51. These results indicate that combining precise segmentation with advanced LLMs like GPT-4o improves diagnostic interpretation. The proposed pipeline is illustrated in **Figure 40**. [63].

>>>26

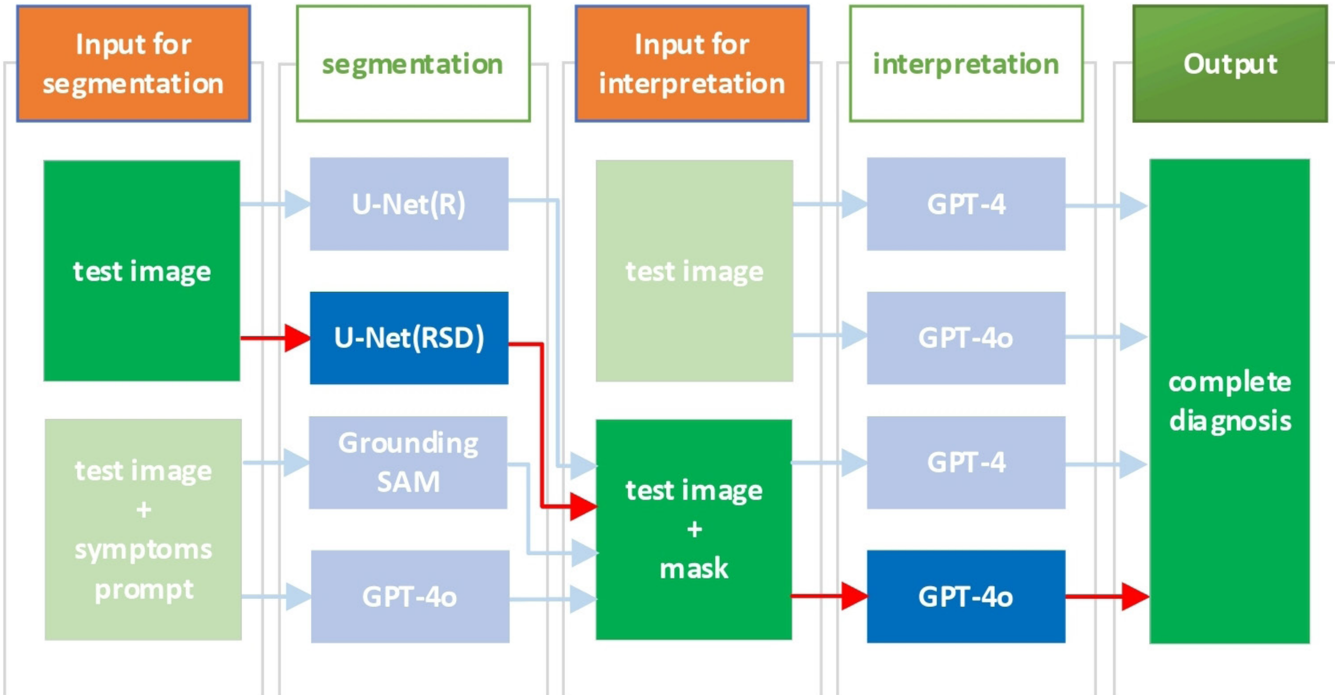

Figure 40. The depicted diagnostic pipeline for canine ophthalmic conditions demonstrates improved performance when both the original images and their corresponding segmented masks from U-Net(RSD) are utilised as inputs for GPT-4o, enhancing the accuracy of the medical diagnoses derived from visual symptoms

The published custom-made DogEyeSeg4 dataset was important in training the models. It provided a diverse set of images reflecting real-world clinical conditions, including variations in lighting, breed, and image quality. The inclusion of synthetically generated images expanded the dataset, allowing models to learn from a broader range of symptom presentations. This comprehensive dataset contributed significantly to the robustness and generalizability of the models.

The limitations identified are the model's sensitivity to image quality, where the diagnostic accuracy is highly dependent on the quality of input images and segmentation masks. Poor image quality or inaccurate segmentation can negatively impact model performance. Also noticed is that models occasionally struggled with conditions like Corneal Ulcer, suggesting the need for further refinement in handling less visually distinct symptoms. Some LLMs, like Llama-3, underperformed in diagnostic interpretation, indicating variability in model capabilities.

Future research should focus on extending the dataset by incorporating more diverse images and medical symptoms, including rare conditions and different breeds, to improve model generalisation. Further usage of advanced model architectures, such as transformer-based models for segmentation to capture global contextual information, could be beneficial. Multimodal data integration by combining image data with clinical history and other modalities could provide more info for a more holistic diagnostic approach. Testing the diagnostic pipeline in clinical settings to validate and refine the models based on practical use should provide valuable feedback.

The experimental results confirm both hypotheses:

Hypothesis 1 is supported by the successful development of a computer vision U-Net(RSD) model for segmentation, and implementation of GPT-4o for interpretation, that accurately recognises canine ocular conditions in images taken in unconstrained environments.

Hypothesis 2 is confirmed by demonstrating that modifications to the U-Net network's input and architecture by using a combination of real and synthetic images, a custom attention mechanism in the form of heatmaps, integrating advanced backbones and employing combined loss functions, lead to better segmentation of canine eye conditions.

6. CONCLUSION

This research has confirmed both initial hypotheses:

Hypothesis 1: By developing a computer vision U-Net(RSD) model for segmentation and implementing GPT-4/GPT-4o for interpretation, we demonstrated that certain canine ocular conditions can be accurately recognised in still images taken in unconstrained environments.

Hypothesis 2: Modifying the U-Net network's input and architecture by using a custom attention mechanism (heatmaps), integrating advanced backbones, and employing combined loss functions resulted in improved segmentation of canine eye conditions.

« BACK TO CONTENTS

These confirmations highlight the effectiveness of advanced deep learning techniques in diagnosing canine eye diseases from images captured in real-world settings. The key contributions of this research are:

An image dataset for machine learning of canine eye diseases: We developed and published the DogEyeSeg4 dataset, which includes a diverse set of real-world clinical images of canine eyes. This dataset reflects variations in lighting, breed, and image quality, providing a robust foundation for training and evaluating machine learning models in unconstrained environments.

Deep convolutional neural network model for recognition of canine eye clinical symptoms and diseases from still images in unconstrained environments: We successfully implemented and tested convolutional neural network models that accurately detect and recognise canine ocular symptoms and diseases in images captured under real-world conditions. The models demonstrated high accuracy in localising the eye region and identifying key ocular symptoms, even in instances such as variable lighting and patient movement.

Deep neural network based on U-Net for segmentation of canine eye clinical symptoms from still images in unconstrained environments: By enhancing the standard U-Net architecture with advanced backbone networks like ResNet34 and employing combined loss functions (Dice and Focal Loss), and augmented input, we developed the U-Net(RSD) model. This model achieved superior performance in segmenting key ocular symptoms, providing the precise delineation necessary for accurate diagnosis.

An improved method for segmentation of canine eye conditions based on U-Net: We introduced modifications to the U-Net network's input and architecture, including synthetized images and the use of a custom attention mechanism in the form of heatmaps generated from SSD eye detection. These enhancements led to improved segmentation accuracy and reduced false positives.

In summary, this research demonstrates the potential of integrating advanced computer vision techniques with deep learning models to enhance the diagnosis of canine ocular diseases. The development of the DogEyeSeg4 dataset and the improved U-Net-based segmentation methods contribute valuable resources and methodologies for future research in veterinary ophthalmology. The successful application of models like U-Net(RSD) and GPT-4o showcases the feasibility of deploying such technologies in real-world clinical settings, ultimately aiming to improve diagnostic accuracy and patient outcomes in veterinary medicine.

Despite these advancements, the research identified limitations such as the models' sensitivity to image quality and difficulties in diagnosing subtle conditions like corneal ulcers. Addressing these challenges is significant for developing more accessible and reliable diagnostic tools. Future research should focus on expanding the dataset, refining model architectures, and integrating multimodal data to further enhance automated veterinary diagnostics and develop more comprehensive approaches to animal health care. V

7. REFERENCES

F. Abdullah *et al.*, "A Review on Glaucoma Disease Detection Using Computerized Techniques," *IEEE Access*, vol. 9, pp. 37311–37333, 2021, doi: 10.1109/ACCESS.2021.3061451.

Y. Hagiwara *et al.*, "Computer-aided diagnosis of glaucoma using fundus images: A review," *Computer Methods and Programs in Biomedicine*, vol. 165, pp. 1–12, Oct. 2018, doi: 10.1016/j.cmpb.2018.07.012.

M. S. Haleem, L. Han, J. van Hemert, and B. Li, "Automatic extraction of retinal features from colour retinal images for glaucoma diagnosis: A review," *Computerized Medical Imaging and Graphics*, vol. 37, no. 7–8, pp. 581–596, Oct. 2013, doi: 10.1016/j.compmedimag.2013.09.005.

L. Li, M. Xu, X. Wang, L. Jiang, and H. Liu, "Attention Based Glaucoma Detection: A Large-Scale Database and CNN Model," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA: IEEE, Jun. 2019, pp. 10563–10572. doi: 10.1109/CVPR.2019.01082.

N. Chakrabarty and S. Chatterjee, "A Novel Approach to Glaucoma Screening using Computer Vision," in *2019 International Conference on Smart Systems and Inventive Technology (ICSSIT)*, Tirunelveli, India: IEEE, Nov. 2019, pp. 881–884. doi: 10.1109/ICSSIT46314.2019.8987803.

Zhuo Zhang *et al.*, "ORIGA-light: An online retinal fundus image database for glaucoma analysis and research," in *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, Buenos Aires: IEEE, Aug. 2010, pp. 3065–3068. doi: 10.1109/IEMBS.2010.5626137.

J. Sivaswamy, S. R. Krishnadas, G. Datt Joshi, M. Jain, and A. U. Syed Tabish, "Drishti-GS: Retinal image dataset for optic nerve head(ONH) segmentation," in *2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI)*, Apr. 2014, pp. 53–56. doi: 10.1109/ISBI.2014.6867807.

S. Kanse and D. Yadav, "Retinal Fundus Image for Glaucoma Detection: A Review and Study," *Journal of Intelligent Systems*, vol. 28, pp. 43–56, Jan. 2019, doi: 10.1515/jisys-2016-0258.

F. Fumero, S. Alayón, J. L. Sanchez, J. Sigut, and M. Gonzalez-Hernandez, "RIM-ONE: An open retinal image database for optic nerve evaluation," presented at the Int. Sym. on CBMS, Jul. 2011, pp. 1–6. doi: 10.1109/CBMS.2011.5999143.

B. Luo, J. Shen, Y. Wang, and M. Pantic, "The iBUG Eye Segmentation Dataset," p. 9 pages, 2019, doi: 10.4230/OASICS.ICCSW.2018.7.

S. J. Garbin, Y. Shen, I. Schuetz, R. Cavin, G. Hughes, and S. S. Talathi, "OpenEDS: Open Eye Dataset," May 17, 2019, *arXiv*: arXiv:1905.03702. Accessed: Nov. 22, 2022. [Online]. Available: http://arxiv.org/abs/1905.03702

H. Proenca, S. Filipe, R. Santos, J. Oliveira, and L. A. Alexandre, "The UBIRIS.v2: A Database of Visible Wavelength Iris Images Captured On-theMove and At-a-Distance," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 8, pp. 1529–1535, Aug. 2010, doi: 10.1109/TPAMI.2009.66.

W. Fuhl, G. Kasneci, and E. Kasneci, "TEyeD: Over 20 million real-world eye images with Pupil, Eyelid, and Iris 2D and 3D Segmentations, 2D and 3D Landmarks, 3D Eyeball, Gaze Vector, and Eye Movement Types," in *2021 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Oct. 2021, pp. 367–375. doi: 10.1109/ismar52148.2021.00053.

A. R. Strom, M. Hässig, T. M. Iburg, and B. M. Spiess, "Epidemiology of canine glaucoma presented to University of Zurich from 1995 to 2009. Part 1: Congenital and primary glaucoma (4 and 123 cases)," *Veterinary Ophthalmology*, vol. 14, no. 2, pp. 121–126, 2011, doi: https://doi.org/10.1111/j.1463-5224.2010.00855.x.

D. A. J. Johnsen, D. J. Maggs, and P. H. Kass, "Evaluation of risk factors for development of secondary glaucoma in dogs: 156 cases (1999–2004)," *Journal of the American Veterinary Medical Association*, vol. 229, no. 8, pp. 1270–1274, 2006, doi: 10.2460/javma.229.8.1270.

M. Boevé and F. Stades, "[Glaucoma in dogs and cats. Review and retrospective evaluation of 421 patients. I. Pathobiological background, classification and breed predisposition]," *Tijdschrift voor diergeneeskunde*, vol. 110, no. 6, p. 219— 227, Mar. 1985.

J. Y. Kim, H. E. Lee, Y. H. Choi, S. J. Lee, and J. S. Jeon, "CNN-based diagnosis models for canine ulcerative keratitis," *Sci Rep*, vol. 9, no. 1, Art. no. 1, Oct. 2019, doi: 10.1038/s41598-019-50437-0.

O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," May 18, 2015, *arXiv*: arXiv:1505.04597. Accessed: Nov. 03, 2022. [Online]. Available: http://arxiv.org/abs/1505.04597

O. Petit, N. Thome, C. Rambour, and L. Soler, "U-Net Transformer: Self and Cross Attention for Medical Image Segmentation," 2021, doi: 10.48550/ARXIV.2103.06104.

R. Azad *et al.*, "Medical Image Segmentation Review: The success of U-Net," 2022, doi: 10.48550/ARXIV.2211.14830.

V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation," Oct. 10, 2016, *arXiv*: arXiv:1511.00561. Accessed: Mar. 02, 2023. [Online]. Available: http://arxiv.org/abs/1511.00561

L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation," in *Computer Vision – ECCV 2018*, vol. 11211, V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, Eds., in Lecture Notes in Computer Science, vol. 11211. , Cham: Springer International Publishing, 2018, pp. 833–851. doi: 10.1007/978-3-030-01234-2_49.

M. Pirhala, J. Goga, V. Kurilova, J. Pavlovicova, and S. Kajan, "Segmentation of Significant Regions in Retinal Images: Perspective of U-Net Network Through a Comparative Approach," in *Systems, Signals and Image Processing*, vol. 1527, G. Rozinaj and R. Vargic, Eds., in Communications in Computer and Information Science, vol. 1527. , Cham: Springer International Publishing, 2022, pp. 29–40. doi: 10.1007/978-3-030-96878-6_3.

P. K. Karn and W. H. Abdulla, "Advancing Ocular Imaging: A Hybrid Attention Mechanism-Based U-Net Model for Precise Segmentation of Sub-Retinal Layers in OCT Images," *Bioengineering*, vol. 11, no. 3, p. 240, Feb. 2024, doi: 10.3390/bioengineering11030240.

N. Siddique, S. Paheding, C. P. Elkin, and V. Devabhaktuni, "U-Net and Its Variants for Medical Image Segmentation: A Review of Theory and Applications," *IEEE Access*, vol. 9, pp. 82031–82057, 2021, doi: 10.1109/ACCESS.2021.3086020.

Z. Liu *et al.*, "Swin Transformer: Hierarchical Vision Transformer using Shifted Windows," 2021, *arXiv*. doi: 10.48550/ARXIV.2103.14030.

E. Xie, W. Wang, Z. Yu, A. Anandkumar, J. M. Alvarez, and P. Luo, "SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers," 2021, *arXiv*. doi: 10.48550/ARXIV.2105.15203.

X. Jia, J. Bartlett, T. Zhang, W. Lu, Z. Qiu, and J. Duan, "U-Net vs Transformer: Is U-Net Outdated in Medical Image Registration?," in *Machine Learning in Medical Imaging*, vol. 13583, C. Lian, X. Cao, I. Rekik, X. Xu, and Z. Cui, Eds., in Lecture Notes in Computer Science, vol. 13583. , Cham: Springer Nature Switzerland, 2022, pp. 151–160. doi: 10.1007/978-3-031-21014-3_16.

K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," Dec. 10, 2015, *arXiv*: arXiv:1512.03385. Accessed: Nov. 03, 2022. [Online]. Available: http://arxiv.org/abs/1512.03385

M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," Sep. 11, 2020, *arXiv*: arXiv:1905.11946. Accessed: Nov. 03, 2022. [Online]. Available: http://arxiv.org/abs/1905.11946

K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for LargeScale Image Recognition," Apr. 10, 2015, *arXiv*: arXiv:1409.1556. Accessed: Nov. 03, 2022. [Online]. Available: http://arxiv.org/abs/1409.1556

C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," Dec. 11, 2015, *arXiv*: arXiv:1512.00567. Accessed: Nov. 03, 2022. [Online]. Available: http://arxiv.org/abs/1512.00567

T. Nemoto *et al.*, "Effects of sample size and data augmentation on U-Net-based automatic segmentation of various organs," *Radiol Phys Technol*, vol. 14, no. 3, pp. 318–327, Sep. 2021, doi: 10.1007/s12194-021-00630-6.

M. Burić, G. Paulin, and M. Ivašić-Kos, "Object Detection Using Synthesized Data," p. 15. In Proceedings of the ICT Innovations Conference, Ohrid, North Macedonia.

A. Dandekar, R. A. M. Zen, and S. Bressan, "A Comparative Study of Synthetic Dataset Generation Techniques," in *Database and Expert Systems Applications*, S. Hartmann, H. Ma, A. Hameurlain, G. Pernul, and R. R. Wagner, Eds., in Lecture Notes in Computer Science. Cham: Springer International Publishing, 2018, pp. 387–395. doi: 10.1007/978-3-319-98812-2_35.

L. Yang *et al.*, "Diffusion Models: A Comprehensive Survey of Methods and Applications," Jun. 23, 2024, *arXiv*: arXiv:2209.00796. Accessed: Jul. 01, 2024. [Online]. Available: http://arxiv.org/abs/2209.00796

R. Po *et al.*, "State of the Art on Diffusion Models for Visual Computing," Oct. 11, 2023, *arXiv*: arXiv:2310.07204. Accessed: Jul. 01, 2024. [Online]. Available: http://arxiv.org/abs/2310.07204

L. Zhang, A. Rao, and M. Agrawala, "Adding Conditional Control to Text-to-Image Diffusion Models," Nov. 26, 2023, *arXiv*: arXiv:2302.05543. Accessed: Jul. 01, 2024. [Online]. Available: http://arxiv.org/abs/2302.05543

N. Carlini *et al.*, "Extracting Training Data from Diffusion Models," 2023, *arXiv*. doi: 10.48550/ARXIV.2301.13188.

A. Stöckl, "Evaluating a Synthetic Image Dataset Generated with Stable Diffusion," 2022, *arXiv*. doi: 10.48550/ARXIV.2211.01777.

C. H. Savage *et al.*, "General-Purpose Large Language Models Versus a Domain-Specific Natural Language Processing Tool for Label Extraction From Chest Radiograph Reports," *American Journal of Roentgenology*, p. AJR.23.30573, Jan. 2024, doi: 10.2214/AJR.23.30573.

A. Q. Jiang *et al.*, "Mistral 7B," 2023, doi: 10.48550/ARXIV.2310.06825.

K. Saab, T. Tu, and W.-H. Weng, "Capabilities of Gemini Models in Medicine," 2024, doi: 10.48550/ARXIV.2404.18416.

W. Huang *et al.*, "How Good Are Low-bit Quantized LLaMA3 Models? An Empirical Study," 2024, doi: 10.48550/ARXIV.2404.14047.

T. Woelfle, J. Hirt, P. Janiaud, L. Kappos, J. P. A. Ioannidis, and L. G. Hemkens, "Benchmarking Human-AI Collaboration for Common Evidence Appraisal Tools," Apr. 22, 2024. doi: 10.1101/2024.04.21.24306137.

T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "BERTScore: Evaluating Text Generation with BERT," 2019, doi: 10.48550/ARXIV.1904.09675.

J. Hessel, A. Holtzman, M. Forbes, R. Le Bras, and Y. Choi, "CLIPScore: A Reference-free Evaluation Metric for Image Captioning," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, 2021, pp. 7514–7528. doi: 10.18653/v1/2021.emnlp-main.595.

K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: a method for automatic evaluation of machine translation," in *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, Philadelphia, Pennsylvania: Association for Computational Linguistics, 2001, p. 311. doi: 10.3115/1073083.1073135.

« BACK TO CONTENTS

M. Denkowski and A. Lavie, "Meteor Universal: Language Specific Translation Evaluation for Any Target Language," in *Proceedings of the Ninth Workshop on Statistical Machine Translation*, Baltimore, Maryland, USA: Association for Computational Linguistics, 2014, pp. 376–380. doi: 10.3115/v1/W14-3348.

K. Ganesan, "ROUGE 2.0: Updated and Improved Measures for Evaluation of Summarization Tasks," 2018, doi: 10.48550/ARXIV.1803.01937.

P. Anderson, B. Fernando, M. Johnson, and S. Gould, "SPICE: Semantic Propositional Image Caption Evaluation," 2016, doi: 10.48550/ARXIV.1607.08822.

M. Burić, M. Ivašić-Kos, and S. Grozdanić, "DogEyeSeg4: Dog Eye Segmentation 4-Class Ophthalmic Disease Dataset." Faculty of Informatics and Digital Technologies, University of Rijeka. Accessed: Aug. 22, 2024. [Online]. Available: https://urn.nsk.hr/urn:nbn:hr:195:405214

S. Grozdanić, S. Đukić, S. Luzhetskiy, N. Milčić-Matić, and T. Lazić, *Atlas bolesti oka pasa i mačaka*. Beograd: Oculus Vet, 2020.

M. Buric, B. Kovacic, and M. Ivasic-Kos, "The Dog Eye Guardian App: From Image to Diagnosis with AI Insights," in *2024 9th International Conference on Smart and Sustainable Technologies (SpliTech)*, Bol and Split, Croatia: IEEE, Jun. 2024, pp. 1–6. doi: 10.23919/SpliTech61897.2024.10612583.

"The EU's General Data Protection Regulation (GDPR) in a Research Context," in *Fundamentals of Clinical Data Science*, Cham: Springer International Publishing, 2019, pp. 55–71. doi: 10.1007/978-3-319-99713-1_5.

M. Buric, S. Grozdanic, and M. Ivasic-Kos, "Diagnosis of ophthalmologic diseases in canines based on images using neural networks for image segmentation," *Heliyon*, p. e38287, Sep. 2024, doi: 10.1016/j.heliyon.2024.e38287.

I. J. Goodfellow *et al.*, "Generative Adversarial Networks," 2014, *arXiv*. doi: 10.48550/ARXIV.1406.2661.

D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," 2013, *arXiv*. doi: 10.48550/ARXIV.1312.6114.

J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," 2020, *arXiv*. doi: 10.48550/ARXIV.2006.11239.

G. Paulin and M. Ivasic-Kos, "Review and analysis of synthetic dataset generation methods and techniques for application in computer vision," *Artif Intell Rev*, vol. 56, no. 9, pp. 9221–9265, Sep. 2023, doi: 10.1007/s10462-022-10358-3.

R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA: IEEE, Jun. 2022. doi: 10.1109/cvpr52688.2022.01042.

J. Thompson, H. Dong, K. Liu, F. He, M. Popescu, and D. Xu, "A Rule-Based Approach for Generating Synthetic Biological Pathways," in *Computational Intelligence Methods for Bioinformatics and Biostatistics*, vol. 13483, D. Chicco, A. Facchiano, E. Tavazzi, E. Longato, M. Vettoretti, A. Bernasconi, S. Avesani, and P. Cazzaniga, Eds., in Lecture Notes in Computer Science, vol. 13483., Cham: Springer International Publishing, 2022, pp. 105–116. doi: 10.1007/9783-031-20837-9_9.

M. Buric and M. Ivasic-Kos, "Next-Generation Computer Vision in Veterinary Medicine: A Study on Canine Ophthalmology," Jul. 22, 2024. doi: 10.36227/techrxiv.172166061.12686289/v1.

E. J. Hu *et al.*, "LoRA: Low-Rank Adaptation of Large Language Models," *arXiv preprint arXiv:2106.09685*, 2023, doi: 10.48550/arXiv.2106.09685.

F. Zhang and M. Pilanci, "Riemannian Preconditioned LoRA for Fine-Tuning Foundation Models," *arXiv preprint arXiv:2403.19776*, 2024, doi: 10.48550/arXiv.2403.19776.

M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting,"

>>>30

in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques - SIGGRAPH '00*, Not Known: ACM Press, 2000, pp. 417– 424. doi: 10.1145/344779.344972.

J. Song, C. Meng, and S. Ermon, "Denoising Diffusion Implicit Models," in *International Conference on Learning Representations*, 2021. [Online]. Available: https://openreview.net/forum?id=St1giarCHLP

P. Luc, C. Couprie, S. Chintala, and J. Verbeek, "Semantic Segmentation using Adversarial Networks," 2016, doi: 10.48550/ARXIV.1611.08408.

M. Buric, M. Ivasic-Kos, and S. Martinčić-Ipšić, "The Disease of the Canine Eye - From Image to Diagnosis Using AI," *WSCG International Conferences in Central Europe on Computer Graphics, Visualization and Computer Vision*, Jun. 2024.

Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D UNet: Learning Dense Volumetric Segmentation from Sparse Annotation," 2016, *arXiv*. doi: 10.48550/ARXIV.1606.06650.

S. Jegou, M. Drozdzal, D. Vazquez, A. Romero, and Y. Bengio, "The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Honolulu, HI, USA: IEEE, Jul. 2017, pp. 1175–1183. doi: 10.1109/CVPRW.2017.156.

V. Dumoulin and F. Visin, "A guide to convolution arithmetic for deep learning," 2016, *arXiv*. doi: 10.48550/ARXIV.1603.07285.

J. Deng, W. Dong, R. Socher, L.-J. Li, Kai Li, and Li Fei-Fei, "ImageNet: A largescale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL: IEEE, Jun. 2009, pp. 248–255. doi: 10.1109/CVPR.2009.5206848.

F. Ramzan *et al.*, "A Deep Learning Approach for Automated Diagnosis and MultiClass Classification of Alzheimer's Disease Stages Using Resting-State fMRI and Residual Neural Networks," *J Med Syst*, vol. 44, no. 2, p. 37, Feb. 2020, doi: 10.1007/s10916-019-1475-2.

M. S. Islam, F. A. Foysal, N. Neehal, E. Karim, and S. A. Hossain, "InceptB: A CNN Based Classification Approach for Recognizing Traditional Bengali Games," *Procedia Computer Science*, vol. 143, pp. 595–602, 2018, doi: 10.1016/j.procs.2018.10.436.

V. Nair and G. E. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines".

A. L. Maas, A. Y. Hannun, A. Y. Ng, and others, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. icml*, Atlanta, GA, 2013, p. 3.

K. He, X. Zhang, S. Ren, and J. Sun, "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification," in *2015 IEEE International Conference on Computer Vision (ICCV)*, Santiago, Chile: IEEE, Dec. 2015, pp. 1026–1034. doi: 10.1109/ICCV.2015.123.

X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *International Conference on Artificial Intelligence and Statistics*, 2010. [Online]. Available: https://api.semanticscholar.org/CorpusID:5575601

C. M. Bishop, *Pattern recognition and machine learning*. in Information science and statistics. New York: Springer, 2006.

M. Buric and M. Ivasic-Kos, "AI in Veterinary Ophthalmology: Canine Eye Diseases Diagnosis with Image Segmentation and Large Language Models," in *World Conference on Smart Trends in Systems, Security and Sustainability - WorldS4*, London, UK: In Press, Jul. 2024.

A. Vaswani *et al.*, "Attention is All you Need," in *Neural Information Processing Systems*, 2017. [Online]. Available: https://api.semanticscholar.org/CorpusID:13756489

T. B. Brown *et al.*, "Language Models are Few-Shot Learners," 2020, *arXiv*. doi: 10.48550/ARXIV.2005.14165.

C. Wang, X. Liu, and D. Song, "Language Models are Open Knowledge Graphs," 2020, *arXiv*. doi: 10.48550/ARXIV.2010.11967.

T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "HighResolution Image Synthesis and Semantic Manipulation with Conditional GANs," 2017, *arXiv*. doi: 10.48550/ARXIV.1711.11585.

M. Hussain, "YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection," *Machines*, vol. 11, no. 7, p. 677, Jun. 2023, doi: 10.3390/machines11070677.

K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," Jan. 24, 2018, *arXiv*: arXiv:1703.06870. Accessed: Mar. 02, 2023. [Online]. Available: http://arxiv.org/abs/1703.06870

S. Liu *et al.*, "Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection," Mar. 20, 2023, *arXiv*: arXiv:2303.05499. Accessed: Apr. 22, 2024. [Online]. Available: http://arxiv.org/abs/2303.05499

A. Kirillov *et al.*, "Segment Anything," Apr. 05, 2023, *arXiv*: arXiv:2304.02643. Accessed: Apr. 22, 2024. [Online]. Available: http://arxiv.org/abs/2304.02643 [90] W. Liu *et al.*, "SSD: Single Shot MultiBox Detector," 2015, doi: 10.48550/ARXIV.1512.02325.

S. Hykes, *Docker, Inc.* (Jan. 05, 2024). [Linux, Windows, macOS]. Available: https://www.docker.com/

G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. TensorFlow Developers, *TensorFlow*.(Jul. 11, 2024). Zenodo. doi: 10.5281/ZENODO.4724125.

G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

A. Clark, "Pillow (PIL Fork) Documentation." readthedocs, 2015. [Online]. Available:https://buildmedia.readthedocs.org/media/pdf/pillow/latest/pillow.pdf

A. Ronacher, *Flask*. (Aug. 05, 2024). Python. [Online]. Available: https://flask.palletsprojects.com/en/3.0.x/

A. Thamizharasan, M. S. Murugan, and S. Parthiban, "Surgical Management of Cherry Eye in a Dog," *Intas Polivet*, 2016.

R. M. Tripathi, D. K. Kashyap, and D. K. Giri, "Surgical Management of Cherry Eye in a Dog," *Intas Polivet*, 2016.

A. F. Siegel, "ANOVA," in *Practical Business Statistics*, Elsevier, 2016, pp. 469– 492. doi: 10.1016/B978-0-12-804250-2.00015-8.

J. W. Tukey, "Comparing Individual Means in the Analysis of Variance," *Biometrics*, vol. 5, no. 2, p. 99, Jun. 1949, doi: 10.2307/3001913.

A. F. M. Alkarkhi and W. A. A. Alqaraghuli, *Easy statistics for food science with R: Abbas F.M. Alkarkhi (Malaysian Institut of Chemical & Bioengineering Technology Universiti Kuala Lumpur, UniKL, MICET, 78000 Melaka, Malaysia),* Wasin A.A. Alqaraghuli (Skill Education Center, PA, A-07-03 Pearl Avenue, *Sungai Chua, 43000 Kajang, Selangor, Malaysia)*. London San Diego Cambridge Kidlington: Academic Press, 2019.

D. C. Montgomery, *Design and analysis of experiments*, Ninth edition. Hoboken, NJ: John Wiley & Sons, Inc, 2017.

SAVA
South African Veterinary Association
Suid-Afrikaanse Veterinêre Vereniging

BRANCH CONGRESSES

join us for a relaxing weekend - catch up with colleagues, get the latest Product updates from trade partners, while earning 10-15 CPD points!

NORTHERN NATAL & MIDLANDS BRANCH

Fordoun Hotel And Spa, Natal Midlands

11-12 OCTOBER 2025

KWA-ZULU NATAL BRANCH

San Lameer Resort, Southbroom

25-26 OCTOBER 2025

SOUTHERN CAPE BRANCH

Oubaai Hotel Golf & Spa, Knysna

31-1 OCTOBER - NOVEMBER 2025

vetlink

For more information visit www.vetlink.co.za

Blindfolding Blind?

Dr. GA Bauer BVScDrMedVet(Ophth)
Gondwana Wildlife Services, Fort Beaufort, Eastern Cape

In Field Slit Lamp Examination of a rhino

Introduction

Chemical immobilisation of wild animals with associated capture and transport has developed from its early days a few decades ago to what it has become today.

Hundreds of animals are chemically immobilised and captured annually in South Africa using various drug combinations, techniques and methods. In many cases, especially in dangerous game, fractious animal species, and most importantly, in white and black rhinoceros, blindfolds are used once the animals are immobilised. This article highlights some cases where application of blindfolds has caused ocular disease.

Blindfolds and their use:

Different materials are used to cover the eyes during immobilisation procedures in game. These types of materials and methods for blindfolds are wide ranging and include *inter alia*:

1. Specifically manufactured halters / head masks as blindfolds,
2. Mutton cloth or similar material wrapped loosely around the head over the eyes,
3. Mutton cloth or similar material wrapped tightly around the head over the eyes,
4. Mutton cloth or similar material wrapped tightly around the head with a "doughnut" ring of tightly wrapped material forming a "standoff" for the mutton cloth off the eye,
5. Items of clothing – shirts, jackets, overalls, etc. draped or tied over the head and eyes,
6. The above materials can either be or not be fixed / held in place by tightly wrapped adhesive tape / insulation tape / Velcro / belt and buckles, etc

Rationale for use of blindfolds

Occlusion of the eyes during chemical immobilisation / capture / transport of animals (viz. rhinoceros) is used for the following reported reasons:

1. Decreasing the effect of visual stimuli on opioid immobilised individuals
2. Safety of staff and people in cases of spontaneous arousals
3. Eliminating the access of direct sunlight to the retina through a mydriatic pupil
4. Calming of partially immobilised individuals
5. Easier handling of partially immobilised individuals by people
6. Prevention of contamination of the eye and its adnexa by foreign material (e.g. dirt, plant matter, etc.)
7. Prevention of desiccation of the cornea
8. Transport of rhinoceros in semi tranquilised state for extended periods of time
9. Protection of the cornea and adnexa in animals immobilised for extended surgical or medical procedures

Clinical cases of ophthalmic disease caused by blindfolds

Over the years of wildlife / veterinary ophthalmology practice, a number of rhinoceros patients have been presented with ocular disease following capture / transport. In all of the cases attended, the animals were transported blindfolded in an anaesthetised / semi immobilised state for a number of hours. In some cases of orphaned calves, the blindfolds were left in place for a few days to facilitate easier handling and nursing to allow calves to latch onto bottles. The lesions observed in these cases ranged from central corneal ulceration to corneal rupture. Progression of the demise of the cornea depended on the infection or not of the open central cornea with collagenase producing agents and / or fungi, the time since initial corneal insult and presentation, and the promptness and efficacy of ocular treatment.

Surgical procedures employed in the treatment of these cases included the following:

1. Temporary tarsorraphy
2. Nictitans flap
3. Conjunctival pedicle grafts
4. Conjunctivocorneal advancement grafts
5. Gundersen Flap
6. Enucleation

A single individual was eventually rendered bilaterally blind, whilst others were rendered with varying degrees of ophathlmic compromise including severe corneal opacity, unilateral enucleation, etc.

Causes of the corneal damage observed

Contrary to what is often believed, corneal damage is not ONLY from direct abrasion of the cornea by the blindfold irrespective of the material used. During immobilisation / partial immobilisation a few physiological changes occur in the eye, including:

1. Decreased tear production
2. Total or partial loss of the blink response
3. Lagophthalmos
4. Reduced spontaneous blinking
5. Decreased corneal sensitivity
6. Total or partial loss of the corneal reflex

Some of the above may be exacerbated physically by tight fitting blindfolds or by blindfolds which have shifted or moved. The above listed physiological changes of the eye, either alone, or coupled with direct contact of any foreign material (blindfold) on the corneal surface, is how these eyes get damaged, sometimes irreversibly.

Alternatives to the use of blindfolds

It is clear, that during any immobilisation / anaesthetic procedure of an animal, protection of the cornea is vital to prevent ocular complications. Although many more animals are captured and exposed to fitment of blindfolds than what have been presented for examination, it is incumbent on us as the primary people responsible for the health of our wildlife patients to explore alternatives in an effort to prevent ocular damage and possibly blindness, either unilaterally or bilaterally in our patients, many of which are the iconic white and black rhinoceros.

« BACK TO CONTENTS

Ophthalmic Protection Alternatives

For short procedures of immobilisation (less than 1 hour) regular topical application of an inert lubricating gel (e.g. KY Jelly) is advocated. For longer periods of immobilisation, topical lubricating gels together with a temporary tarsorraphy is advocated.

If performed correctly, the latter procedure protects the cornea and conjunctiva and there is no possibility of contact with the cornea or conjunctiva by foreign materials.

Failure to remove the tarsorrpahy sutures prior to reversal / release of the animal is the worst possible complication in this procedure.

Images of a few cases observed:

Central Corneal Erosion

Central Corneal Erosion / Neovascularisation

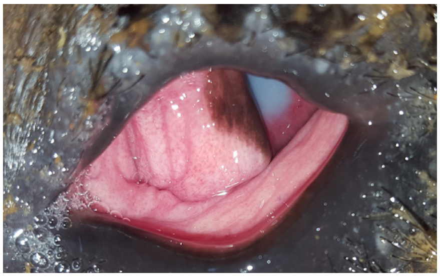

Severe Conjunctivokeratitis

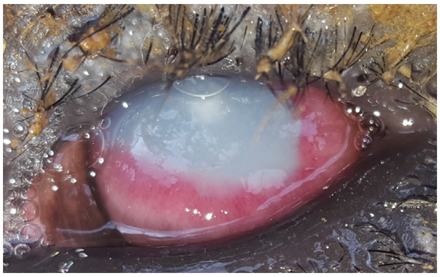

Severe bullous Keratopathy with Neovascularisation

Ruptured corneal Ulcer

Ruptured Corneal Ulcer

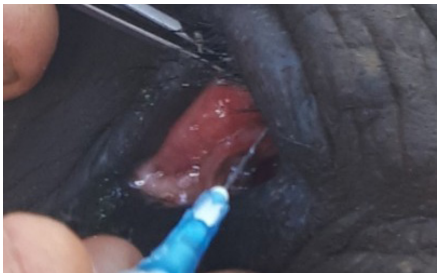

Bulbal Subconjunctival Injection

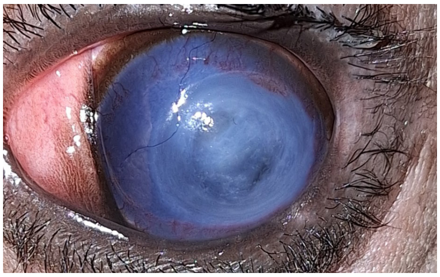

Conservative therapy outcome

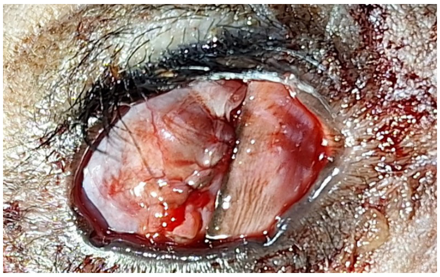

Gundersen flap

SAVETCON named PCO of the Year at the 2025 Women in MICE Awards!

We are thrilled to announce that SAVETCON has been awarded the **Professional Conference Organiser (PCO) of the Year** at the 2025 Women in MICE Awards! This national accolade is a powerful reflection of our team's unwavering commitment to delivering events that are not only exceptional, but also innovative and sustainable.

Redefining Conferences and Events

At SAVETCON, we believe in reimagining what events can be. Over the years, we've embraced new technologies, fresh ideas, and sustainable practices to create meaningful, impactful experiences. From introducing smarter digital solutions that cut down on waste to partnering with responsible vendors, sustainability has become part of who we are. It's this combination of creativity and responsibility that resonated with the judges and set us apart.

Why This Recognition Matters

The Women in MICE Awards honour the women driving transformation in the **M**eetings, **I**ncentives, **C**onferences, and **E**xhibitions sector. They shine a spotlight on leadership, innovation, and the bold ideas shaping the future of the industry. Being named PCO of the Year is more than an award for us - it's a recognition of our passion, our clients' trust, and our belief in raising the standard for professional events.

A Collective Achievement

This award belongs to our incredible team. Their passion, creativity, and tireless dedication are what make SAVETCON a leader in the industry. Each project they deliver reflects not only their skill, but also their drive to push boundaries and craft experiences that leave a lasting impact.

Gratitude to Our Clients and Partners

We also extend a heartfelt thank you to our clients and partners. Your trust and collaboration are at the heart of this achievement. Every successful event is built on these partnerships, and we are grateful to continue creating remarkable experiences with you.

Looking Ahead

Winning PCO of the Year motivates us to keep innovating, to keep championing sustainability, and to keep creating unforgettable moments. The future of events is evolving - and so are we. From new technologies to eco-conscious strategies, SAVETCON is ready to set new benchmarks for the industry in South Africa and across our borders!

In addition, Corné Engelbrecht was awarded the **Lifetime Achievement Award for 2025** at the prestigious Awards Ceremony.

"Winning the Women in MICE Lifetime Achievement Award is an incredibly special milestone for me, both personally and professionally.

On a personal level, it's such an affirmation of the passion, resilience, (often sacrifices) and dedication I've poured into my work over the biggest part of 29 years. Much of it has been behind the scenes, but always with a clear purpose: to make a difference and leave an impact. In the Event industry, every day is a mix of coffee, chaos, and creating unforgettable moments, and I wouldn't trade it for anything. It reminds me that every challenge, every late night, every blister and all the tears, frustration and every moment of creative problem-solving has been worth it.

Professionally, this recognition gives me a platform to inspire and empower others in the Event industry - especially women who are still finding their space, voice and carving out their own paths. It strengthens my commitment to excellence and fuels my passion for mentorship, motivating me to keep raising the bar. Most importantly, this award isn't just about me - it's about the incredible mentors, colleagues and clients who have been part of my journey. Without them, this would not have been possible."

Thank you for being part of our journey - this is only the beginning!

South African Veterinary Association
Suid-Afrikaanse Veterinêre Vereniging

Notice is hereby given that the 120th Annual General Meeting of the South African Veterinary Association will be held on

Saturday, 22 November 2025

08:00

at

VetHouse
47 Gemsbok Avenue, Monument Park, Pretoria or via Zoom

Please notify the secretary, Ms Sonja Ludik, whether you will be attending in-person or via Zoom

Join Zoom Meeting

https://us06web.zoom.us/j/82563504181?pwd=8tTblw0ixPDwQi2JeX2Vbn15oqstbq.1

Meeting ID: 825 6350 4181
Passcode: 083426

Tel: 012 346 1150
sonja@sava.co.za

By order of the Board
Registered office:
47 Gemsbok Ave, Monument Park, 0181

August 2025

CREDO

We, the members of the Association, resolve at all times:

- To honour our profession and the Veterinary Oath
- To maintain and uphold high professional and scientific standards
- To use our professional knowledge, skills and resources to protect and promote the health and welfare of animals and humans
- To further the status and image of the veterinarian and to foster and enrich veterinary science
- To promote the interests of our Association and fellowship amongst its members.

Ons, die lede van die Vereniging, onderneem om te alle tye:

- Ons professie in ere te hou en die Eed na te kom
- 'n Hoë professionele en wetenskaplike peil te handhaaf en te onderhou
- Ons professionele kennis, vaardigheid en hulpbronne aan te wend ter beskerming en bevordering van die gesondheid en welsyn van dier en mens
- Die status en beeld van die veearts te bevorder en die veeartsenykunde te verryk
- Die belange van ons Vereniging en die genootskap tussen sy lede te bevorder.

Welcome to Shaye Hughes

I am delighted to introduce myself to the readers of VETNews.

My name is Shaye Hughes, a married mom to a busy and bright 5-year-old boy who keeps me on my toes and reminds me daily of the importance of balance between family and career.

For the past 16 years, I have worked in the financial industry, where I have built extensive experience in accounting, tax, compliance, and business advisory services. Over the years, I have had the privilege of working closely with individuals, small businesses, and organisations across different sectors, helping them navigate the often-complex world of finance with confidence.

Currently, I am part of the dynamic team at Robeli, an accounting firm. At Robeli, we pride ourselves on being more than just accountants – we are trusted partners to our clients. Our services range from tax and compliance to payroll and company secretarial support. We are especially passionate about working with clients in niche industries, such as wildlife, veterinary practices, and non-profits, where tailored financial guidance can make a meaningful difference.

Being part of Robeli has allowed me to combine my love for numbers with my passion for helping people and businesses achieve growth and sustainability. I look forward to engaging with the veterinary community and contributing to conversations that strengthen both financial and professional success in this vital industry.

Shaye

Dear SAVA member

Please note that Vethouse will be closed from 12:00pm on Tuesday 23 December 2025 and will reopen on Monday, 05 January 2026.

SAVA wishes all its members and their families a joyful festive season and everything of the best for the year ahead.

Kind regards
Sonja Ludik
Interim General Manager

Geagte SAVV lid

Neem asseblief kennis dat Vethuis gesluit sal wees vanaf 12:00pm op Dinsdag 23 Desember 2025 en weer op Maandag, 05 Januarie 2026, sal heropen.

Die SAVV wens alle lede en hul families 'n vreugdevolle feestyd, en alles van die beste vir die nuwe jaar toe.

Vriendelike groete
Sonja Ludik
Waarnemende Algemene Bestuurder

Have you joined yet?

Our WhatsApp Communication Platform is now active!
Stay updated on member news and important topics in the veterinary industry.

Joining the Community: A Step-by-Step Guide

Prerequisites:

- Ensure you have WhatsApp installed on your device.
- Please add/save SAVA's WhatsApp number to your contacts: 081 849 6088.
- To join, you will need an invitation link or an invite from a community admin or an existing SAVA member.

Joining via an Invitation Link:

1. Obtain the invitation link from a SAVA member or community admin.
2. Click the link, which will open WhatsApp and prompt you to JOIN.
3. Tap "Join" to confirm your membership.

We look forward to connecting with our members on this community platform!

Did not receive the link invite? Are you encountering other issues?

Please contact Sonja van Rooyen at +27 (0) 12 346 1150 or at assistant@sava.co.za

VACCINATION BOOKS

The LAST ORDERS for 2025 for vaccination books will be on 30 November.

No orders will be accepted after this – this is to ensure that stock will be dispatched by 15 December 2025.

Orders for 2026 will resume on 15 January.

To order / for more information contact Debbie Breeze on 012 346 1150 or debbie@sava.co.za


SAVA
South African Veterinary Association

"...to use our professional knowledge, skills and resources to protect and promote the welfare of animals and humans.."

Mission

The South African Veterinary Association aims to serve its members and to further the status and image of the veterinarian. We are committed to upholding the highest professional and scientific standards, and to utilising the professional knowledge, skill and resources of our members, to foster close ties with thecommunity and thus promote the health and welfare of animals and mankind.

Tel:+27(12)3461150/1, www.sava.co.za, vethouse@sava.co.za

SAVA Community Veterinary Clinics (SAVA-CVC)

A project of theSAVA which aims to provide primary healthcare to pets owned by owners who cannot afford the services of a veterinarian. Various clinics have been established countrywide. Veterinarians donate their time and skill. Financial support is required, primarily to purchase medication and materials.

Registered non-profit company(1998/016654/08) non-profit organisation (000-234NPO) and public benefit organisation(130001321)

Tel:+27(12)3461150/1, www.communityvet.co.za, cvc@sava.co.za

South African Veterinary Foundation (SAVF)

The foundation was established to promote a greater understanding of animals, through promoting research and an informed public. Student bursaries and research grants are awarded; research results are published in national and international journals.

www.savf.org.za, savf@sava.co.za

South African Veterinary Council (SAVC)

The SAVC is the regulatory body for the veterinary profession and veterinary para-professions in South Africa.

Tel:+27(12)3456347,www.savc.org.za,savc@savc.org.za

State Veterinary Service

Importandexportrequirements

www.daff.gov.za → Branches → Agricultural Production, Health and Food Safety → AnimalHealth → Import/Export Policy Unit → (application forms can be found here)

Build client loyalty

Donate to the Pet Memorial Fund

SOUTH AFRICAN VETERINARY FOUNDATION
SUID-AFRIKAANSE VETERINÊRE STIGTING
www.savf.org.za

In Memoriam

Dr Clive Marwick

On Friday evening, 25 July 2025, we bade farewell to a giant in veterinary medicine. A man whose life's work shaped not only an industry but countless lives. Dr Clive Marwick was more than a practitioner-he was a pioneer, mentor, visionary, and friend.

Clive's legacy as a rural, large animal practitioner in the Eastern Free State is etched into the pastures and hearts he served. Known as the "father of bovine herd health in SA," he upheld a gold standard of professionalism, compassion, and integrity. His approach transcended treatment-he championed prevention, education, and collaboration, bringing healing not only to animals but to the communities that depended on them.

Messages of condolence have arrived from around the globe, from colleagues and protégés who called him a mentor and described his impact as unforgettable. Their words echo the heart of Clive's legacy: generous in knowledge, patient in teaching, unwavering in support.

Clive het in verskeie nasionale en internasional professionele organisasies met buitengewone integriteit gedien. Om maar net 'n paar van die organisasies te noem, die Suid Afrikaanse Veterinêre Vereniging, die Suid Afrikaanse Veterinêre Raad, die Onderstepoort Veterinêre Fakulteitsraad, die World Veterinary Association, en vele meer. Clive was deurentyd deur sy kollegas en andere genomineer om op finansiële komitees van organisasies en instellings te dien.

Daar het hy díe organisasies en instellings uitgedaag om die regte ding te doen; en hulle finansiële sake in orde te kry. Onder sy leierskap is vooruitgang gemaak ten beste van toekomstige generasies.

In his leadership of the South African Veterinary Council, Clive transformed vision into reality. As President, he led decisively- with merit, fairness, and faith. His belief in inclusivity, sound decision-making, and always doing the right thing wasn't just spoken-it was lived. He guided the Council to purchase its own building, strengthened its finances, and helped forge it into a highly effective institution. As Registrar of the SA Veterinary Council, I knew and worked with Clive for more than twenty years, and I saw firsthand the ease with which he led- wise, humble, sincere and deeply committed to the greater good.

Clive faced challenges with humour and optimism. He saw possibility in every problem and light in every dark corner. He made people feel seen, heard, and valued. Above all, he lived his life as a man of God- with quiet dignity, grace, and purpose.

His legacy is a living one—found in healthier herds, stronger institutions, and not least of all, the hearts of those he mentored.

Clive wou sien dat sy professie met dieselfde professionalisme, integriteit, geloof, opregtheid en deernis voortuitgaan.

We salute you Clive! Ons groet jou, Clive, mag jy jou Hemelse reis in Vrede voortsit V

Lynette Havinga

In Memoriam

Dr Paul Bartels

It is with profound sadness that we announce the passing of Paul Bartels, a dedicated veterinarian and conservationist, esteemed colleague, and beloved friend, who left us on 3rd July 2025. Paul was a remarkable figure in the field of conservation, known for his passion, innovative ideas, and unwavering commitment to wildlife conservation.

Born on the 22nd February 1956, Paul's journey as a scientist and veterinarian started at the University of Stellenbosch, where he earned a BSc Degree in 1979. He completed his training as a veterinarian at the Faculty of Veterinary Science in 1986, where his passion for wildlife and assisted reproduction was evident, and obtained an MSc degree from the Mammal Research Institute of the University of Pretoria in 1996.

His career was marked by significant contributions to wildlife conservation, biobanking, reproductive science and veterinary science. He was an affiliate of the University of Johannesburg, being awarded an Associate Professor post. He worked tirelessly at various institutions, including the National Zoo, SANBI, and the Smithsonian, and was a collaborator on groundbreaking projects at the Wildlife Breeding Centre (WBRC) that he established at Pelindaba.

He was appointed to the Tshwane University of Technology as a Senior Lecturer and, latterly, to a research affiliate post with Stellenbosch University.

Throughout his career, Paul was well-known within numerous conservation hierarchies – overseas zoos such as Toronto Zoo and Australia, and actively engaged with the IUCN in contributing to the establishment of genome resource banks and innovative biobanking techniques.

His research on double-headed spermatozoa and cloning in cattle in 2003 positioned him as a pioneer in South Africa's conservation efforts.

Paul's legacy is evident in his numerous publications and his role as an excellent marketing manager and fundraiser. He was instrumental in empowering young conservationists through his work with WESSA, where he served as chairman and board member, promoting community-led conservation and anti-poaching initiatives.

His passion extended beyond his professional life; Paul was always ready to lend a helping hand and share his knowledge with others. He was personable and jovial, building strong relationships with students and colleagues alike. His grand ideas and innovative approaches, such as lab-grown meat and the Nyoka Vulture restaurant, enriched the lives of many.

Paul's legacy includes the establishment of the SA Biobank Network and his involvement in the Leopard Project and Magaliesberg conservation efforts. He was also a microlight pilot, often seen flying over his farm, where he created an airstrip for the community.

He is survived by his partner, Elize Venter and son, and will be deeply missed by all who had the privilege of knowing him.

Paul Bartels was more than a colleague; he was a mentor, a visionary, and a friend. His passion for life and conservation will continue to inspire future generations. V

Richard Burroughs

CALL FOR ABSTRACTS

GUIDELINES:

The abstract must be a complete but concise synopsis of the content that will be both informative and useful to the potential audience. The format for the abstracts is flexible, but should contain a clear objective, methods and results (where applicable) and a logical conclusion.

Two independent people will review abstracts and successful authors will be notified within two weeks following the deadline fro submission. Authors must accept the terms and conditions for speakers. Abstracts that do not comply with the following format will be returned to the authors for corrections:

- Must be in English
- Abstracts may not exceed 350 words.
- Must clearly indicate the author's name. Write names of authors in the order of last names and asterisk (*) the person who will be the presenter.
- Must clearly indicate the organization (name and location of institution, clinic, practice etc). If the affiliation of co-workers differs, indicate the names and affiliations by corresponding numbers (Superscripts).
- Must contain a list of references.

Deadline: 30 November 2025

https://vetlink.co.za/saeva-congress-2026/

Influential Life Coaching

THE CURRENCY

Dr Mats Abatzidis

B.Sc. B.V.Sc.

New Insights Certified VIP Life Coach
mats.abatzidis@yahoo.co.za
Founder of Influential Life Coaching
http://www.matsaba.wix.com/drmatscoach
Author of the published book "Life outside your comfort zone. Better and beyond all expectations".
http://www.amazon.com/s/ref=nb_sb_noss?url=search-alias%3Ddigital-text&field-keywords=Abatzidis
http://www.life-coach-directory.co.za/mats-abatzidis

I hope this article finds you and your families well and safe!!! Last month, we looked at shifting one's state of mind from a negative one to a positive one within minutes. In this instalment, we are going to look at the currency in each of our lives. Namely, TIME, MONEY and KNOWLEDGE.

We have all heard the famous quotes:

- "Time is money"
- "Knowledge is power"
- "Knowledge has a beginning, but no end"
- "Time is the most valuable thing a man can spend"
- "The more you learn, the more you earn"

Then I came across another quotation that got my attention: **"There are three different currencies in life; TIME, MONEY and KNOWLEDGE, whichever you want abundantly, you use the other two to get it".** This is the first time I have seen someone mention all three of these entities (time, knowledge and money) in one statement. This rule is quite simple and applicable to anyone.

I made my decision to become a veterinarian when I was eight years old, because of my own dog being injured in a dog fight, and me not having the knowledge to help him feel better. I then spent the next 18 years gaining as much knowledge as I could to succeed at graduating as a registered veterinarian. For me to gain the necessary **knowledge**, I had to spend **time and money**. We have all been there, where we paid for learning a skill, a profession, or a trade. The latter applies to any form of education, whether it is formal (through an institute) or informal (books, the internet, reading articles, etc). Sometimes it will be required that we pay with our time (practising, speaking with a mentor, etc.) or with our money, and sometimes it requires both. In my experience, the following equation applied:

TIME + MONEY = KNOWLEDGE

I then qualified from Onderstepoort, and I was ready to start working as a registered veterinarian. That meant that it was time to **apply** all **the knowledge** I had acquired through my studying years and apply those skills daily **for several hours** a day so that I could earn a salary. I was planning on getting married, buying a house, and buying a car at the very least. The funds for all of those "wants" had to come from somewhere. There are multiple ways **to earn money**. Having a job is one of them. Starting my own practice and becoming self-employed would have been an alternative. So, when we apply the above to the equation, it shifts to look as follows:

KNOWLEDGE + TIME = MONEY

I had reached the part of my life where I could **exchange time for money** while using my skills. The next step would be to take a break from time to time. We are not robots, and even if we were, we would still need to rest intermittently.

That means that we work and **apply our knowledge to earn money**, which also allows us **to then have** some **time** off to rest, travel and eventually retire. As expected, I am certain you have seen the pattern by now. The three variables in the equation shift again, and the calculation becomes as follows:

KNOWLEDGE + MONEY = TIME

This is all great and well, but inevitably, there are disclaimers in life. Perhaps you have guessed it already! All there of these currencies or entities are not the same. **The one that is NOT REPLENISHABLE is time!** "Lost time is never found again. – *Benjamin Franklin*. That is to say that as valuable as time is as an asset, it is limited and once it has passed, it is quite rare that we get it back. Time should not be wasted. We all have 24 hours in a day.

The **decision that needs to be made**, and a decision that should never be taken lightly, is how you want to use yours. I am not only referring to time taken off for resting or while you are on holiday. I am talking about your daily routines and your basic use of time on a regular basis.

Do we make a conscious decision to make **money** or to acquire **knowledge**? Robert Kiyosaki mentions this in his books - "...the difference between a rich person and a poor person is how they use their time..."

THINK ABOUT IT RIGHT NOW

You're reading this article, which means you're currently using your **time** and probably your internet connectivity, which you've bought with your **money**. If you decide that you have not gained anything from this instalment, you will have wasted your **time and money**. That's how it works.

Next month, we will continue looking at more ideas for improving our quality of life and overall performance, both at work and at home. V

World Sight Day

Dr Anthony Goodhead
Specialist Veterinary Ophthalmologist
Johannesburg and Cape Animal Eye Hospitals
www.animaleyehospital.co.za

World Sight Day is observed every year on the second Thursday of October.

During the 'SightFirst' campaign in 2000, the Lions Clubs International and the International Agency for the Prevention of Blindness (I.A.P.B.) proclaimed every second Thursday of October to be observed as World Sight Day.

The primary objective was to draw the public's attention to the importance of following safe practices to prevent blindness and other problems related to vision in people.

This day reminds us to take care of our eyes and to show love and support to those who can't see. This concept is not only reserved for people but can certainly be extrapolated to our furry friends.

The number of eye-related cases in general practice is rather high. It is estimated that about 3 % of all veterinary consultations in animals are eye-related, and on average, it is estimated that 68% of cases actually require referral. In the case of horses, it may be about 7% of consultations that involve the eye. In this segment on ophthalmology, we shall highlight a few simple things that we should be considering when it comes to routine pet eye care health.

1. **Many clients request a routine "eye check".** These should be encouraged, just like any regular health check for dental disease, lameness, etc. It is amazing to see how often, during a full ocular examination, further pathology is found that the client was unaware of. Often, this pathology may have a profound effect on the eye or even involve regular medication or surgical intervention. The most common problems that we incidentally encounter include:

a. Distichiae or ectopic cilia.
Often these small eyelid cilia may make contact with the corneal surface and result in discomfort, blepharospasm, keratitis or even ulceration.

These cilia are frequently noted during a slitlamp biomicroscopy examination, and epilation is done if the dog is undergoing surgery for some other eye-related problem.

b. Eyelid neoplasia
Tarsal margin masses seldom get smaller, so it is not ideal to tell an owner to watch the lesion and re-assess later. It is always easier and more successful to resect a small amount of eyelid margin than to battle with closing a larger tarsal margin wound that results in a poorer cosmetic result. Once a mass has reached a third or more of the length of the eyelid margin, then some form of eyelid splitting or skin flap is required to close the eyelid to allow sufficient blinking. The golden rule is to remove masses when they are small.

c. Mature Cataracts
Or developing cataracts are frequently seen as incidental observations, unbeknownst to the client, and further monitoring may be valid for such cases. Cataract surgery should be considered when there is a diffuse lens opacity, or even if vision problems have been observed. Probably the most important reason to perform regular eye care examinations in cataract patients is to constantly look out for any signs of lens-induced uveitis. This is the inflammatory reaction that occurs as a result of the antigenic lens proteins leaking through the lens capsule. This is a very serious and common problem in diabetic patients.

Clinically, the first signs may be a slight blepharospasm, lacrimation or a congested bulbar conjunctiva. Closer examination with a focal light source in a darkened room may demonstrate aqueous flare or subtle oedema to the cornea. These patients require urgent medical management with topical corticosteroid or NSAID eye drops, or even oral NSAIDs. The globe should then be assessed for cataract surgery by performing gonioscopy, ocular ultrasound and electroretinography.

The goal is to remove the antigenic lens proteins by phacoemulsification cataract surgery as soon as the globe has quietened down.

d. Retinal degeneration.
Signs of retinal degeneration are usually noted on a fundus examination, where hyper-reflectivity of the tapetal region is noted, together with atrophy of retinal blood vessels. Often, the clients are not aware of any signs of vision loss at home until an obstacle course is performed. We perform this test in a room with a dimmer switch, and it often becomes obvious to the client that as we reduce the light intensity, the dog becomes less active and more likely to knock into obstacles placed in the room. Often, only then that the client mentions situations at home where their pet may have a vision problem. Often, this test alone reassures them that their pet is managing, and they can then do similar testing at home and manage their pet's environment better as the vision progressively gets worse. Although there is no effective treatment for retinal degeneration, many clients want this problem monitored.

2. **Regular periorbital grooming** is necessary in many breeds with long fur around the face. The show-ring grooming practices are often impractical [such as a long fringe in an Old English Sheepdog] or the classic Schnauzer cut. There simply is no reason to have long hair hanging into a dog's eyes, as this causes chronic irritation, lacrimation and even secondary ulcerations.

Having said that, it is very important that if one is trimming facial hair on facial folds with scissors or a clipper, such as in Shih Tzu and Pekingese dogs, then the facial fold hair should be made very short and maintained in this position. If this cut border, then grows out over time, it may result in a "brush" border of fur, irritating the cornea and causing damage or chronic pigmentary keratitis to develop.

« BACK TO CONTENTS

3. Tear staining: It is always important to try and establish the reason for tear staining, which, for many owners, is offensive. In such cases, the problem may be caused by chronic irritation such as entropion, distichiae, ectopic cilia, eyelid neoplasia, agenesis of lacrimal puncta, blocked tear ducts, etc. The brown staining seen is a porphyrin pigment and clinically is not abnormal. Pigment is thought to become more intensely dark in colour when exposed to light. The chronic use of pet wipes may assist. Historically, people have used long-term oral antibiotics, especially Tylosin. This concept is contrary to all antibiotic stewardship issues and should not be condoned.

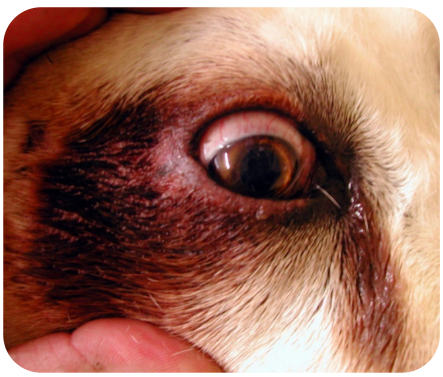

4. The brachiocephalic breeds such as Shih Tzu, Pekingese, Pugs, etc, often develop a medial pigmentary keratitis in very young animals. It is not uncommon that pigmentation is often seen at about 4-5 months of age in these breeds. During the early routine consultations and vaccination examinations, one can look out for the pets that have a very large palpebral fissure opening. They have the "bulging eyes". They make ideal candidates for prophylactic treatment. Probably the most profound prophylactic surgery to prevent medial pigmentary keratitis for such patients is a bilateral medial canthoplasty.

This procedure assists in closing down the medial canthus, and in so doing, the medial entropion and caruncle hair are removed, and the palpebral fissure is made smaller, preventing any facial fold irritation.

This surgery does require magnification and the use of fine Vicryl sutures.

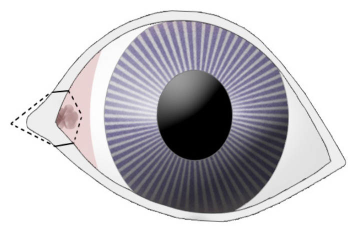

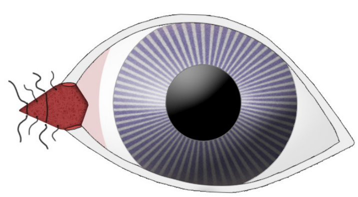

5. During a regular eye care examination, the presence of a mucoid ocular discharge around the eye is a warning bell for potential keratoconjunctivitis sicca [Dry eye]. This is rarely a bacterial infection, contrary to what clients may be thinking. In such cases, a Schirmer Tear Test is essential to confirm the diagnosis. Dry eye is the most common eye problem seen by medical and veterinary ophthalmologists.

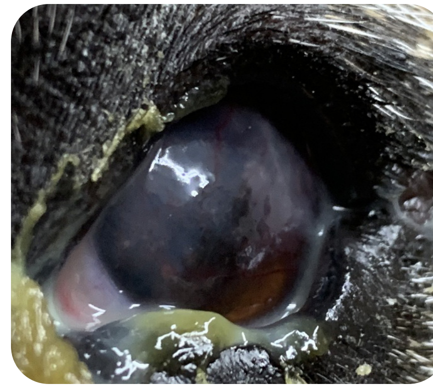

Fortunately, most cases fall into the idiopathic / Immune-mediated group and respond favourably to Tacrolimus 0.02% in Oil. Such cases do require regular supervision as signs of pigmentary keratitis, vascular keratitis and scarring may require further intervention in the future.

6. "Red Eye" - One of the most useful reasons for having a regular eye care consult, or at least a good ocular examination when presenting pets for other reasons, is to take note of a red eye. This is frequently noted, and on closer examination, one may find very important signs that could suggest bigger problems. Any congested globe would be an early warning sign for conditions such as uveitis, glaucoma or Dry Eye.

7. National Hereditary Eye Scheme – For breeders of pedigree dogs, this scheme, which has been operational in South Africa for 30 years, allows breeders to present dogs for examination with the view to looking for hereditary eye conditions.

This is an effective method to screen dogs and identify dogs that are affected with known, or presumed to be hereditary, eye disease before they are used for breeding. There is an official certificate that is issued after such an examination. By paying attention to the above simple eye health concepts, one can increase the standard of veterinary care offered to our patients. In this way, we can turn World Sight Day into an everyday occurrence. V

Eye COLUMN

More than a Gut Feeling: Clinical Insights and Nutritional Management of Diarrhoea and Constipation in Dogs and Cats

By Janine Volschenk BSc BVSc (UP)

Gastrointestinal (GI) disorders are a daily reality in companion animal practice, and diarrhoea and constipation remain two of the most common clinical signs presented in both dogs and cats. They represent opposite ends of the spectrum of abnormal stool consistency and frequency, yet both significantly impact patient health, client concern, and quality of life. While often seen as routine, these conditions can be complex, multifactorial, and challenging to manage. Alongside diagnostics and medical management, nutrition is increasingly recognised as a cornerstone of therapy.

Diarrhoea: When Things Move Too Fast

Diarrhoea is broadly defined as an increase in stool water content, often accompanied by increased frequency and volume. Clinically, it may be acute (sudden onset, often self-limiting) or chronic (lasting more than three weeks, requiring more thorough investigation).

Common causes include:

- Infectious agents (e.g., Parvovirus, Giardia, Coccidia, Tritrichomonas foetus in cats)
- Dietary indiscretion or abrupt diet change
- Inflammatory bowel disease (IBD) and food-responsive enteropathies
- Exocrine pancreatic insufficiency (EPI)
- Neoplasia or systemic disease (e.g., hepatic, renal, endocrine disorders)

Clinical approach:

A good history and faecal scoring system helps differentiate small-bowel from large-bowel diarrhoea. Diagnostics often include faecal analysis, haematology/biochemistry, imaging, and sometimes endoscopy with biopsy. Supportive treatment may involve fluid therapy, antiemetics, probiotics, and in some cases antimicrobials—but only when indicated.

Nutritional management:

Diet plays a crucial role in both acute and chronic diarrhoea. Highly digestible diets help reduce intestinal workload, while fibre blends help normalise stool quality. Soluble fibres (e.g., psyllium) ferment in the colon to produce short-chain fatty acids (SCFAs) such as butyrate, which nourish enterocytes and support mucosal health. Insoluble fibres (eg, cellulose) add bulk and help regulate transit time.

Hydrolysed protein diets are invaluable in managing food-responsive enteropathies, while therapeutic GI diets with tailored fibre blends are designed to stabilise stool quality in both dogs and cats.

Constipation: When Things Slow Down

Constipation is characterised by infrequent, difficult, or absent defecation. While more commonly encountered in cats—particularly older or obese individuals—it can also affect dogs. Chronic cases may progress to obstipation or megacolon, especially in felines.

Common causes include:

- Mechanical obstruction (e.g., pelvic stenosis, masses, strictures)
- Painful defecation (e.g., orthopaedic disease, anal sac disease)
- Neurological disorders (e.g., spinal cord disease, nerve injury)
- Dehydration or electrolyte imbalances (e.g., hypokalaemia, hypercalcaemia)
- Idiopathic megacolon in cats
- Environmental/behavioural factors (e.g., stress, lack of exercise, or inadequate litter box access)

Clinical approach:

Work-up typically includes abdominal palpation, imaging (radiographs and/or ultrasound), and assessment for underlying systemic disease. Mild cases may respond to rehydration, enemas, or laxatives, but recurrent constipation requires a long-term plan.

Nutritional management:

Diet is pivotal in chronic cases. Fibre can both prevent and manage constipation, but the type and balance matter. Soluble fibres help retain water in the stool, softening faeces, while insoluble fibres increase bulk and stimulate peristalsis.

« SCAN THIS QR CODE TO VIEW THE CANINE AND FELINE FAECAL SCORING CHART (ROYAL CANIN)

Some cats with megacolon benefit from high-fibre diets, while others do better on highly digestible, low-residue diets that minimise stool volume. Tailoring fibre levels to the individual patient is key. Adequate water intake is equally important—wet diets or water fountains may help in cats.

Nutrition as a Therapeutic Cornerstone

Whether stools are too loose or too firm, nutrition plays a fundamental role in restoring balance. Fibre is not just "roughage"—it is a dynamic tool that influences transit time, stool quality, microbial composition, and mucosal health.

In practice:

- **Acute diarrhoea** →Easily digestible, moderate fibre diets support recovery.
- **Chronic diarrhoea or IBD** →Hydrolysed or limited antigen diets with tailored fibre help reduce antigenic load and improve gut health.
- **Constipation** →Fibre adjustment (either increase or decrease) and water optimisation are first-line strategies, with diet forming the foundation of management.

Fibre Type	Physiological Effect	Examples	Clinical Relevance
Soluble fibre	Absorbs water, forms gels; slows transit	Psyllium	Improves stool consistency in diarrhoea; nourishes colonocytes
Insoluble fibre	Adds bulk, increases stool volume, stimulates motility	Cellulose	Useful for constipation by promoting peristalsis
Fermentable fibre	Supports microbiota, produces SCFAs, reduces inflammation	FOS, MOS, beet pulp	Enhances gut health, modulates immune function
Non-fermentable fibre	Increases faecal bulk without fermentation	Cellulose, Psyllium	Beneficial in constipation, especially megacolon

Table 1. Summary of the different fibre types and their clinical relevance

Modern therapeutic diets, such as those formulated for gastrointestinal support, incorporate a balance of soluble and insoluble fibres, prebiotics, and highly digestible ingredients to provide clinicians with powerful tools in everyday GI cases.

Final Thoughts

Diarrhoea and constipation may be opposite clinical signs, but both highlight the central role of the gut in overall health. While diagnostics and medical management remain critical, nutrition is often the piece that makes the difference between temporary relief and successful long-term management. Understanding how to leverage dietary strategies empowers veterinarians and nurses to improve outcomes for both patients and owners and enhance quality of life.

References

1. Hall EJ, German AJ. *Diseases of the Small Intestine.* In: Ettinger SJ, Feldman EC, eds. Textbook of Veterinary Internal Medicine, 8th ed. Elsevier, 2017.

2. Washabau RJ, Day MJ, Willard MD, et al. *WSAVA International Gastrointestinal Standardization Group Guidelines on Small Animal Gastrointestinal Diseases.* J Small Anim Pract. 2010;51(6):384–392.

3. Jergens AE, Heilmann RM. *Canine Chronic Enteropathy—Current State-of-the-Art and Future Directions.* J Vet Intern Med. 2014;28(3):689–698.

4. Marks SL, Laflamme DP, McCluggage DG. *Management of Constipation in Cats.* Compend Contin Educ Vet. 2011;33(7):E1–E6.

5. Royal Canin Veterinary Diets. Gastrointestinal and Fibre Response Diet Formulations. [Product Information].

BVH
Bryanston Veterinary Hospital
&
Specialist Centre

SPECIALIST REFERRAL HOSPITAL

- Open 24 Hours
- General and Referral Practice
- Emergency and Critical-care Facility
- Overnight Hospitalisation with Veterinary supervision

- Telephone (011) 706-6023 (All Hours)
6 Ballyclare Drive, Bryanston

email: bvh@global.co.za
web: bryanstonvet.co.za

SPECIALIST REFERRALS

Surgery: Dr Marthinus Hartman
Surgery: Dr Ross Elliott
Surgery: Dr Klaas-Jan van de Wetering
Medicine: Dr Ryan Friedlein
Anaesthesia & Critical Care: Dr Gareth Zeiler
Anaesthesia: Dr Etienne Basson
Ophthalmology: Dr Lo-An Odayar
Theriogenology : Dr Susan Fouché
Exotic Pets & Birds: Dr Dorianne Elliott

012 991 3573
referral@valleyfarmvet.co.za
What's App 066 302 2901

Valley Farm Animal Hospital

JHB ANIMAL EYE HOSPITAL
CAPE ANIMAL EYE HOSPITAL

Johannesburg
Tel: 011 465 1237
Fax: 011 465 2916
44 Kingfisher Drive,
Fourways

Cape Town
Tel: 021 930 6632
Fax: 021 939 9941
1 Uys Krige Drive,
Panorama

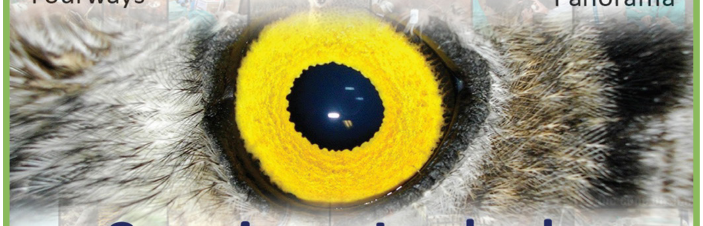

Services Include

- Phaco emulsification for cataracts with intra ocular lens placements.
- Amnion grafts for the treatment of corneal ulcers and sequestrums.
- Placement of suprachoroidal cyclosporine implants for the treatment of equine recurrent uveitis.
- Cysplatin bead implants for the treatments of eyelid, third eyelid and conjunctival squamous cell carcinomas.

WWW.ANIMALEYEHOSPITAL.CO.ZA
info@animaleyehospital.co.za

The SAVA Stress Management Hotline

The following SAVA members are available on the SAVA stress management hotline. If required, they will refer you to professionals.

Ken Pettey: 082 882 7356
ken.pettey@gmail.com

Aileen Pypers : 072 599 8737
aileen.vet@gmail.com

Willem Schultheiss :
082 323 7019
schultheisswillem@gmail.com

Mike Lowry: 084 581 2624
mikelowry@sai.co.za

Tod Collins: 083 350 1662,
tcollins@isat.co.za

VETSCAPE
REFERRALS - EQUINE

We have an equine veterinarian position available at Vetscape, located in the Western Cape, South Africa

This position would be well-suited to someone with a special interest in sport horses but a high level of involvement in all aspects of our hospital will be required.

If you would like to know more about the role or to apply for the position, please email Brittany Reinhardt - brittanyr@vetscape.co.za

Please direct any queries to Brittany Reinhardt at either the above email address or telephonically on +2721 867 0700

Classified Advertisements
Snuffeladvertensies

VETERINARIAN / VEEARTS

WESTERN CAPE / WES-KAAP

GARDEN ROUTE

Vacancy Location:
Garden Route SPCA
Job title: Veterinarian
Job Type: Full-Time
Degree in Veterinary Science
(BVSc or equivalent)
Registered with SAVC
We offer a competitive salary with 08:00 – 17:00 weekday working hours only.
Closing date is 15 October 2025
CVs to be sent to:
chericooke@grspca.co.za
Ref25OC01

FOR SALE / TE KOOP

ANAESTHETIC MACHINE

A new veterinary anaesthetic machine with MSS3 vaporiser or TEC4 refurbished vaporiser. All machine assembly, servicing, and calibrations continue to be applied per the Blueprint of the former Ebrahim Cassim, Chief Anaesthetic Technician of Groote Schuur Hospital. For further information, call Cassim Veterinary Anaesthetics on 021 705 2880 / 082 681 9742.
Email: encassim@gmail.com or visit www.cvanaesthetics.co.za
Ref13JA01

VARIOUS ITEMS FOR SALE

From the closure of a practice, all at more than half the original cost and all in good condition
(1) Mindray DP 50 Veterinary Ultrasound Unit with a Mindray Mico Convex Transducer R24 900 (¼ of the original cost)
(2) Olympus CX 23 Microscope R8 400
(3) Mindray Pulse Oximeter R2 600
(4) Edan Pulse Oximeter R3 300
(5) Barcode Scanner R550
(6) Humphrey SEVOFLUORANE Gas Anaesthetic Machine and Emergency trolley R17 774 (¼ of the original cost)
(7) Idexx Catalyst One Analyser R44 100 (¼ of the original cost)
(8) SA 232 Autoclave Steriliser 16 litre R15 400
(9) Refractometer R1 275
(10) Bonart Ultrasonic Dental Descaler with Pack of Dental instruments R5 360
(11) Carestream Vita LE CR System with one 24* 30 Cassette and one 30*40 Cassette R43 725 (¼ of the original cost)
(12) Stainless Steel Cradle R550
(13) Instrument Trolley R1 250
(14) Andis Clipper and blades R2000
(15) Lead Apron R1 050, and Mittens R1 900
Please send a WhatsApp to 0823382678 or an email to everton1@iafrica.com if you are interested in any of the above
Ref25AU01

GENERAL / ALGEMEEN

RADIATION ONCOLOGY

(Referral Practice)

Dr Georgina Crewe BVSc, MSc(Wits)
Radiation therapy may be used alone or in conjunction with surgery and chemotherapy. Radiation is particularly useful in the treatment of solar-induced Squamous cell carcinoma, cutaneous mast cell tumours and sarcomas. Palliative radiation is successful for most tumours as the tumour shrinks and the peripheral nerves are released, relieving the pain caused by the tumour. For more information or to discuss a case, please contact: Georgina Crewe,
115 9th Ave, Fairland, Johannesburg
Telephone: 011-678-3121
Cell: 082-492-6247
E-mail: crewegeorgina@gmail.com
Ref25MY02.

ADVERTISE IN VETNEWS MAGAZINE

Tel:
012 346 1150
Email:
assistant@sava.co.za

24-Hour
Toll-Free
Helpline
0800 21 21 21

Wildlife Group of the SAVA

SAVE THE DATE

Wildlife Group of the SAVA

CONGRESS 2026

26° South Hotel , Muldersdrift, 12-14 March 2026

SAVA
South African Veterinary Association
Suid-Afrikaanse Veterinêre Vereniging

CALL FOR ABSTRACTS

- Must be in English
- Abstracts may not exceed 350 words.
- Must clearly indicate the author's name. Write names of authors in the order of last names and asterisk (*) the person who will be the presenter.
- Must clearly indicate the organization (name and location of institution, clinic, practice etc). If the affiliation of co-workers differs, indicate the names and affiliations by corresponding numbers (Superscripts).
- Must contain a list of references.

Deadline: 30 November 2025

26° south
bush boho hotel

26° South Hotel offers all the comforts required for an unforgettable stay, within the soothing serenity of the African bush. The Magaliesberg mountain range forms a picturesque backdrop for weddings, corporate getaways and leisure seekers alike. Situated within the Cradle Moon Conservancy in Muldersdrift Gauteng, this private reserve provides the opportunity to escape the bustle of the city but is a convenient 20 minute drive from Lanseria International Airport and 1 hour drive from OR Tambo.

https://vetlink.co.za/wildlife-group-2026/

RUVASA

Ruminant Veterinary Association of South Africa

Save the Date

CONGRESS 2026

Champagne Sports Resort, Drakensberg, KwaZulu Natal

8-10 June 2026

THE NEW NORMAL

ADAPTING RUMINANT VETERINARY PRACTICE TO EMERGING BIOSECURITY DEMANDS – STRATEGIES FOR SUSTAINABILITY AND SUCCESS

CALL FOR ABSTRACTS

GUIDELINES:

The abstract must be a complete but concise synopsis of the content that will be both informative and useful to the potential audience. The format for the abstracts is flexible, but should contain a clear objective, methods and results (where applicable) and a logical conclusion.

Two independent people will review abstracts and successful authors will be notified within two weeks following the deadline fro submission. Authors must accept the terms and conditions for speakers. Abstracts that do not comply with the following format will be returned to the authors for corrections:

- Must be in English
- Abstracts may not exceed 350 words.
- Must clearly indicate the author's name. Write names of authors in the order of last names and asterisk (*) the person who will be the presenter.
- Must clearly indicate the organization (name and location of institution, clinic, practice etc). If the affiliation of co-workers differs, indicate the names and affiliations by corresponding numbers (Superscripts).
- Must contain a list of references.

Deadline: 30 November 2025

CHAMPAGNE SPORTS RESORT

CENTRAL DRAKENSBERG

Overlooking the Drakensberg mountain range, this polished resort, set in an elegant building with a thatched roof, is 6 km from the Drakensberg Canopy Tour.

Featuring mountain views, the warm rooms come with TVs, ceiling fans, and tea and coffeemakers. Plush suites add separate lounges, fireplaces and balconies. There are also 1- to 3-bedroom chalets with kitchens, living rooms and dining areas.

Amenities include a buffet restaurant and a sports bar. There are also 5 tennis courts, 5 outdoor pools and a golf course, plus a gym, a spa and event space. Helicopter rides and bike rentals are also available.

https://vetlink.co.za/ruvasa-congress-2026/