

This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

AnimalFormer: Multimodal Vision Framework for Behavior-based Precision Livestock Farming

Ahmed Qazi Taha Razzaq Asim Iqbal* Tibbling Technologies

asim@tibbtech.com

Abstract

We introduce a multimodal vision framework for precision livestock farming, harnessing the power of GroundingDINO, HQSAM, and ViTPose models. This integrated suite enables comprehensive behavioral analytics from video data without invasive animal tagging. GroundingDINO generates accurate bounding boxes around livestock, while HQSAM segments individual animals within these boxes. ViTPose estimates key body points, facilitating posture and movement analysis. Demonstrated on a sheep dataset with grazing, running, sitting, standing, and walking activities, our framework extracts invaluable insights: activity and grazing patterns, interaction dynamics, and detailed postural evaluations. Applicable across species and video resolutions, this framework revolutionizes noninvasive livestock monitoring for activity detection, counting, health assessments, and posture analyses. It empowers data-driven farm management, optimizing animal welfare and productivity through AI-powered behavioral understanding.

1. Introduction

The fusion of computer vision and deep learning has catalyzed a paradigm shift in non-contact animal monitoring systems. These advanced systems, now a cornerstone in modern agriculture, are pivotal for animal behavior quantification and early disease detection [10]. As farming moves towards precision livestock farming, the role of such technologies in ensuring animal health and welfare has become indispensable. Within this domain, pose estimation [13, 23, 29] and instance segmentation [13, 32] emerge as two pillars foundational to behavioral analysis, driving efforts towards creating robust and efficient models. Pose estimation, a domain thoroughly explored, benefits from extensive datasets [33, 35] that provide a wealth of annotated information across diverse animal species. Leveraging such resources, generalizable pose estimation models [31, 34] achieve remarkable results, enhancing our understanding of animal behavior across the animal kingdom. While pose estimation offers in-depth behavioral insights, instance segmentation opens doors to analytics that transcend posture analysis [32]. However, the isolated application of these techniques only provides a fraction of the potential insight.

Current approaches for evaluating the behavioral states of farm animals, such as direct observation or physiological measurements, are often disruptive, subjective, and impractical for large-scale application [6]. There's a lack of scientifically validated metrics to accurately gauge farm animal well-being using these methods [28]. Standards in animal welfare predominantly highlight negative indicators, like harm and stress, rather than positive ones [25, 30]. Yet, recognizing and promoting positive states are crucial not only for the animals' health but also for enhancing farm productivity and product quality. Non-invasive monitoring offers a promising solution to these challenges. By leveraging technologies that minimize stress and bias, we can gather reliable data on animal welfare, potentially improving both the lives of livestock and the outputs they yield [26, 27]. The importance of such advancements cannot be overstated. With the global population projected to reach nearly 10 billion by 2050 [21], the demand for sustainable and efficient farming practices is at an all-time high.

We introduce a revolutionary non-contact, multimodal AI framework, AnimalFormer, that synergizes state-of-theart models for a holistic analysis of livestock behavior. Utilizing an open-source sheep dataset [18], we demonstrate comprehensive behavioral analytics across various activities such as grazing, running, and sitting/resting. The combination of ViTPose [31], GroundingDINO [20], and HQ-SAM [17] allows for nuanced pose estimation and precise segmentation, capturing the granularity of animal movements and interactions. AnimalFormer provides a scalable solution that can be adapted across different livestock species and varying farm sizes, offering insights that are critical for optimizing feed efficiency, monitoring health, and potentially improving breeding programs. This frame-

^{*}Corresponding author.



Figure 1. Our integrated analysis framework, designed for comprehensive behavioral understanding of sheep in a dataset. The framework combines ViTPose and Grounding DINO for pose estimation and contextual understanding respectively, with the high-quality instance segmentation capabilities of HQ-SAM. By fusing these components, our pipeline provides precise keypoints and segmentation masks, essential for in-depth ethological studies. This block diagram provides a clear overview of the data flow and processing steps within our end-to-end solution.

work transcends the capabilities of existing methodologies, offering a novel approach to farm management that is both advanced and operationally efficient. By harnessing raw video data, our framework as shown in **Figure 1**, facilitates precision livestock farming and proactive health monitoring without the need for invasive tagging or interaction. The insights gleaned are not only pivotal for individual animal welfare but also hold the promise of transforming farming practices on a global scale. The non-invasive nature of our technology ensures minimal stress to animals, leading to more accurate behavioral observations and, consequently, a deeper understanding of animal well-being.

To the best of our knowledge, such a comprehensive, multimodal AI framework using latest state-of-the-art Vision Transformers (ViTs) [8] models has not been previously devised specifically for the agriculture and farming community. By bridging this gap, we not only contribute to the advancement of precision livestock farming but also set a new benchmark for animal welfare and farm management practices. The potential for scaling this framework to incorporate additional behavioral analytics and apply it to a wider range of animal species further emphasizes its significance and the broad impact it could have on the future of farming and agricultural industry.

2. Related work

The advancement in non-contact animal monitoring systems, particularly through pose estimation, has significantly improved our understanding of animal behavior. Largescale open-sourced datasets, combined with deep learning techniques, have fostered the development of efficient pose estimation models [34]. These models enable detailed behavior analysis across various species, offering insights into their natural patterns and interactions. Notably, mice have been a focal point in behavior analytics, leading to models that capture their diverse behaviors [12, 23]. In livestock farming, pose estimation has proven invaluable. It facilitates behavior analytics for cattle, aiding in breeding reference, early detection of lameness [2], and monitoring overall health [36]. This approach, free from human intervention, not only ensures comprehensive behavior analytics but also reduces the risk of injury to animals. However, pose estimation alone may not capture the full spectrum of animal interactions, especially in scenarios involving occlusion or requiring the analysis of intricate behaviors such as eating patterns. This limitation stems from the reliance of pose estimation models on keypoints, which might not effectively capture changes across consecutive frames. To address these gaps, instance segmentation emerges as a powerful tool, enabling the detailed tracking of animal interactions and behaviors [32]. Coupled with multiple-object detection, this approach enhances the tracking and trajectory analysis of animals within video footage.

Recognizing the strengths and limitations of both pose estimation and instance segmentation, we propose an endto-end multimodal framework that integrates these methodologies. This framework aims to provide comprehensive behavioral analytics, encompassing a broader range of activities and interactions. Through this synthesis, we seek to offer a more holistic view of animal behavior, bridging the gap between individual and collective behaviors.

3. Methods

3.1. Dataset

The dataset [18] used in this paper was curated for sheep behavior analysis. It constitutes of an extensive collection of videos capturing five distinct sheep activities: *grazing*, *running*, *sitting*/resting, *standing*, and *walking*. The videos were recorded using high-resolution digital cameras from varying angles and positions, and the entire dataset comprises of 417 videos. Although, the duration of each video varies, the total footage is 59 minutes long, totaling to 149, 327 frames. For our analysis, we only consider the videos pertaining to grazing, running and sitting/resting activities. Furthermore, to enhance computational efficiency, we downsampled the dataset by selecting every 10^{th} frame. This downsampling process reduces the computational load while preserving essential information for analysis, ensuring the dataset remains suitable for developing and testing computer vision algorithms aimed at non-invasive monitoring and analysis of sheep behavior.

The activities captured in the dataset are reflective of typical behaviors exhibited by sheep in a natural setting, making it a valuable resource for training deep learning models to recognize and classify these activities accurately. The dataset's format, consisting of high-resolution videos and their corresponding frames (with a notable omission of frames containing humans for privacy), allows for a wide array of computer vision applications, including but not limited to, activity recognition and behavior analysis.

3.2. AnimalFormer

In AnimalFormer, we integrate ViTPose [31], Grounding DINO [20], and HQ-SAM [17] to extract keypoints and perform instance segmentation. These elements collectively contribute insights into the sheep dataset. The comprehensive framework of AnimalFormer is shown in Figure 1.

3.2.1 Pose estimation

We utilize ViTPose [31], the current state-of-the-art pose estimation model, to extract keypoints for each video. ViT-Pose provides a simple but efficient ViT baseline for pose estimation in animals as well as human datasets. It uses plain and hierarchical vision transformers, as the backbone model which is trained on ImageNet-1K [7] with masked image modelling (MIM) [15] pre-training. ViT is coupled with light weight decoders which bypass the use of complex mechanisms including skip-connections and crossattentions and instead comprise of simple deconvolution and prediction layers.

An input image $X \in \mathcal{R}^{H \times W \times 3}$, is first converted into tokens via the patch embedding layer i.e $F \in \mathcal{R}^{\frac{H}{d} \times \frac{W}{d} \times C}$ where *d* is the downsampling ratio for each patch embedding layer and *C* is the number of channels in the input image. The embedded tokens are then fed to the transformer layer which converts them to feature i.e $F_{out} \in \mathcal{R}^{\frac{H}{d} \times \frac{W}{d} \times C}$. Each transformer layer consists of multi-head self attention layer (MHSA) along with a feed-forward network (FFN) i.e.

$$F'_{i+1} = F_i + \text{MHSA}(\text{LN}(F_i)) \tag{1}$$

$$F_{i+1} = F'_{i+1} + \text{FFN}(\text{LN}(F'_{i+1}))$$
(2)

where *i* is the output of the *i*th transformer layer and the initial feature (F_0) is the output of the patch embedding layer. LN is the normalization layer.



Figure 2. Qualitative outputs of our framework depicting various behaviors of sheep. Top row: Running frames with keypoints and bounding boxes illustrating movement dynamics. Middle row: Grazing frames showcasing sheep engaged in feeding with extracted poses and segmentation masks highlighting the focus areas. Bottom row: Resting frames with segmentation masks delineating individual sheep in a state of repose. Each behavior is analyzed through a combination of visual features extracted from images.

The features extracted by the ViT are fed to the decoders which upsamples them twice before converting them to keypoints $K \in \mathcal{R}^{\frac{H}{4} \times \frac{W}{4} \times N_k}$ where N_k is the number of keypoints to be estimated, which is set to 17 for animal datasets.

3.2.2 Animal detection

For our animal detection module, GroundingDINO [20] is utilized, building on the DINO framework's self-supervised learning strengths. DINO, underpinned by the distilled knowledge self-attention [4, 16] mechanism from ViTs, empowers our model to recognize textual descriptions within visual inputs efficiently. The DINO methodology incorporates self-supervised learning strategies [5, 11, 14], utilizing an exponential moving average of the student model's parameters to update the teacher model, thus preventing collapse without a direct contrastive loss [5].

The distilled knowledge within the student network is refined through temperature-scaled softmax functions, as expressed below:

$$P_s(x)^i = \frac{\exp(g_{\theta_s}(x)^i / \tau_s)}{\sum_{q=1}^Q \exp(g_{\theta_s}(x)^q / \tau_s)},$$
(3)

where $P_s(x)^i$ denotes the probability predicted by the student network for the i^{th} class, $g_{\theta_s}(x)^i$ is the logit or the

raw output from the student network for the i^{th} class before applying softmax, τ_s represents the temperature parameter influencing the softmax function's sharpness and the sum in the denominator runs over all Q possible classes, ensuring that the probabilities sum to 1. A similar equation pertains to the teacher network P_t with temperature τ_t .

Grounding DINO advances this architecture by precisely linking textual phrases with their corresponding objects in images, for instance, identifying a cat and a table from an input image and associating them with their respective labels extracted from the input text. This enhanced capability is achieved through an image backbone for extracting visual features, a text backbone for text feature extraction, a feature enhancer for cross-modality fusion, and a crossmodality decoder for refining object boxes.

3.2.3 Animal segmentation

Our segmentation utilizes HQ-SAM [17], an adept extension of SAM [19], for its proficiency in generating highquality masks. The model introduces the *HQ-Output Token*, improving dynamic convolution through a dedicated threelayer MLP (MLP_{dyn}) and an attention mechanism (Attn) that refines the token with input features (F_{in}):

$$HQ-Mask = MLP_{dyn}(HQ-Output Token) \odot F_{HQ-F}$$
, (4)

HQ-Output Token_{und} = Attn(HQ-Output Token,
$$F_{in}$$
), (5)

with the enhanced output feature (F_{out}) attained by element-wise multiplication of the softmax-weighted input feature map:

$$\mathbf{F}_{\text{out}} = \operatorname{softmax} \left(\mathbf{W}_{f} \cdot \mathbf{F}_{\text{in}} \right) \odot \mathbf{F}_{\text{in}}, \tag{6}$$

Global-local fusion combines contextual and boundary detail features across network stages, culminating in a comprehensive feature map (F_{HQ-Fs}):

$$F_{HQ-Fs} = Conv \left(F_{early} \oplus F_{final} \oplus F_{mask} \right), \tag{7}$$

HQ-SAM thus upholds SAM's zero-shot segmentation efficacy with substantial quality improvements for practical, high-fidelity segmentation tasks.

3.3. Behavior analytics

3.3.1 Running videos

Gait. In order to quantify the overall gait pattern of an animal, we look at the varying poses it exhibited during running. A set of 17 keypoints is extracted from each frame of an animal's running video as shown in **Figure 2**. The keypoints are converted to meaningful embedding by applying Uniform Manifold Approximation and Projection (UMAP) [24], which represent the different poses demonstrated by

an individual sheep during the course of its running. For UMAP, we set n_neighbors=20 to encapsulate all the poses emphasized by the running of an animal and a two-dimensional embedding space (n_components=2) is used for visualization. To quantify the unique poses present across the videos and extract similar gait patterns between different sheeps, we cluster the embeddings via K-Means Clustering [22]. We create 10 unique clusters for all the poses demonstrated by the sheep's in the videos. In order to extract meaningful insights about the sheep behavior we assess how gait patterns affect speed, by conducting a speed analysis of the running videos. We opt for UMAP as part of our pipeline due to its ability to preserve local and global structures within the data along with being highly efficient in processing large volumes [24].

Speed. We assess the speed of individual subjects based on their observable movement across video frames (Figure 2) in running videos. Given a set of masks corresponding to each subject's position in sequential frames, we commence by identifying the largest mask within each frame, which we associate with the primary subject of interest. This is predicated on the notion that the primary subject occupies the largest area within the field of view. To ascertain the centroid of the largest mask, we calculate the geometric center of its binary mask area, which serves as a proxy for the subject's position. The temporal evolution of the centroid between consecutive frames allows us to deduce the raw speed of the subject. This is achieved by measuring the displacement of the centroid over time, considering the frame rate of the video and the number of frames skipped between measurements to enhance accuracy. However, raw speed alone does not provide a complete picture due to potential variations in the subject's distance from the camera, which can skew perceived speed. To mitigate this, we normalize the speeds by the areas of the masks. The normalization process involves adjusting the speeds based on a reference area value derived from the mean of all mask areas. The intent is to correct for perspective distortion, wherein objects that are closer to the camera appear to be moving faster than those further away. By normalizing the speeds, we aim to achieve a more accurate representation of the actual speed of the subjects. Our methodology yields a set of normalized speeds for each subject, corresponding to their movement across the frames. This data along with the gait patterns forms the basis for subsequent analyses of the sheep.

3.3.2 Grazing videos

To perform analysis of the grazing behavior among sheep, we begin with segmenting individual sheep across video frames as shown in (**Figure 2**). After accurate identification and segmentation, we determine the keypoint representative of the sheep's nose. This keypoint serves as the focal point



Figure 3. Behavior Analytics. **A.** UMAP representation of the unique gait patterns of the sheep extracted from their running videos. **B.** Speed profile of the animal extracted at the commencement, midpoint, and conclusion of their running videos. **C.** The unique clusters of existing sheep gait patterns. **D.** Spread of different patterns within a single animal across different clusters. **E.** The grazing activity of the sheep in herds vs single. **F.** UMAP representation of the unique resting pattern of the animals in herd vs single.

around which we establish an approximate grazing area. We further refine this selected grazing patch by subtracting the sheep's segmentation masks to exclude non-grazing elements such as sheep's head, ensuring that only relevant grazing area is measured. Within this isolated region, we quantify the green signal intensity. The measurement of the green signal intensity across consecutive frames provides an approximation of the grazing activity which was indicative of the amount of grass intake. This allowed us to make inferential analyses about the grazing patterns among the sheep.

3.3.3 Resting videos

On sheep resting behavior, we employed an image-based analysis to differentiate between individual and herd resting states from front and side view (**Figure 2**). The methodol-

ogy involved segmenting video footage into frames, using bounding boxes to identify and isolate sheep within each frame. We extracted masks for each sheep identified by the bounding boxes, categorizing the data into four distinct groups - front view herd, side view herd, front view single, and side view single. The masks were resized to a standard dimensions for consistency across the dataset. The processed data was then analyzed using UMAP for dimensionality reduction, allowing us to visualize and cluster different resting states. UMAP's configurations were set to n_neighbors=50 to account for the overarching grouping tendencies of sheep, and min_dist=0.01 to achieve distinct clusters indicative of individual and group resting states. A two-dimensional space projection with n_components=2 was chosen for ease of visualization and analysis, with the Euclidean metric due to its natural fit for the data's geometric properties. To quantify the unique resting poses, we cluster the UMAP embeddings via K-Means Clustering. We create 10 unique clusters for all the poses demonstrated by the sheep's in the sitting videos. This approach facilitated the identification of patterns in sheep resting behavior, such as preferences for resting alone or within a group.

4. Results & Discussion

4.1. Indirect relationship between animals' gait diversity and running speed

Since sheep are prey animals with extremely heightened senses of flight and fright, they tend to run away from humans and other animals [9]. Running is a common behavioral habit observed in sheep which can tell us a great deal about their gaits and speed. Using the running videos of individual sheep and their corresponding UMAP embeddings, we extract the different poses exhibited by the sheep, as shown in Figure 3 (A). Each unique color in Figure 3 (A) represents an individual animal's poses. We observe that although for majority animals, the poses exhibited lie closer to each other however in some instances they are spread out. To validate this observation, we cluster the animal poses into 10 unique groups as shown in Figure 3 (C). Although the running patterns of a sheep are independent of other subjects, we observe that some sheep share gait patterns as they belong to the same cluster. Figure 3 (D) shows the different poses encompassed within an individual animals running pattern along with the distribution of poses in multiple clusters. Even though different sheep have overlapping gait patterns, however Figure 3 (D) clearly shows that an individual sheep's gait pattern is consistent within itself and is largely found within a single cluster. Animal IDs 025 and 011, for example, have a high number of poses, which are found in a single cluster i.e cluster ID 9 and 1 respectively. We observe a similar trend in case of a dynamic gait pattern

spanning multiple clusters and observe that majority of the poses are present in a specific cluster. Animal ID 014, for example, has poses spanning clusters 0, 2 and 6 however majority of its poses are present in cluster 6 while a few are distributed among 0 and 2. A similar trend is observed for other animals with higher number of unique poses, including animal IDs 029 and 033, for which majority of their poses are present in cluster 2 and 0 respectively.

Although, pose estimation provides insights regarding the gait patterns, however in order to quantify the animal's health, we further provide speed analysis. Figure 3 (B) illustrates the speed trends of individual sheep at various intervals - the commencement, midpoint, and conclusion of their running sessions. The figure presents a compelling narrative; sheep that display a more limited range of dynamic gait patterns tend to achieve higher speeds. Animal ID 025, for example, encompasses few poses within its running pattern, however it is the fastest among the lot. On the other hand, animal ID 014, is among the slower running sheep despite having a higher number of poses. This phenomenon suggests a possible trade-off between speed and the variety of movements within the gait. In essence, sheep that adhere to a consistent gait pattern, without significant variations, generally exhibit higher velocities. Conversely, sheep with a broader array of dynamic poses show a tendency toward slower speeds. The observed inverse relationship aligns with the principles of locomotive efficiency, where streamlined, repetitive movements facilitate faster motion [1]. This insight is not only of academic interest but also carries practical implications. Understanding the gait diversity and corresponding speed could aid in identifying potential health issues, optimizing shepherding practices, and enhancing overall herd management.

4.2. Enhanced grazing activity in isolation

Through animal detection and pose estimation, we quantify and compare the grazing patterns of sheep in an isolated setting with when they are flocked together. Our results, as shown in Figure 3 (E), reveal compelling evidence that sheep engage in more substantial grazing activity when isolated as compared to when they are part of a herd. In absence of social stimuli and potential threats, the sheep focus more intently on feeding and have a higher grazing activity. This behavior is attributed to the diminished distractions and the lack of need for social interaction or vigilance, which is more prevalent in group settings. In controlled environments, devoid of predators or environmental stressors, animals typically exhibit a reduction in the vigilance behavior that characterizes wild settings. Our observations are in alignment with this principle, indicating that solitary grazing sheep capitalize on the safety and tranquility to maximize their food intake. We observe that when the animals flock together, they tend to consume lesser food, since they pause their grazing for various social interactions, such as establishing social hierarchies or bonding. The quantitative data gleaned from our analysis provides a clear contrast between the two scenarios. Sheep in solitude not only graze more but also show consistent feeding patterns without significant interruptions. This contrasts with herd scenarios where the grazing patterns are interspersed with nonfeeding activities. Our study offers valuable insights into the grazing patterns of domestic sheep, which can be instrumental in enhancing farm management practices and optimizing feeding strategies to ensure the well-being and productivity of the animals.

In short, our examination of sheep grazing behaviors yielded an intriguing pattern - in controlled scenarios, sheep displayed an increased grazing activity when alone compared to being in a herd. This increase is hypothesized to be due to the reduced distractions and social obligations that typically interrupt feeding in a group environment. Our data suggests that in such isolated snippets, the absence of broader social dynamics allows sheep to feed more efficiently, leveraging the security and stillness of a controlled setting to focus on grazing. It is worth noting, however, that these findings are based on brief observational periods. A long-term comprehensive study could reveal different grazing dynamics, as sheep might exhibit varying grazing intensities over extended periods [3]. Our insights are constrained by the scope of the dataset, which did not capture extended durations or include variables like environmental shifts and detailed social interactions, limiting the breadth of our analysis.

4.3. Resting pose similarities across animals

In Figure 3 (F), the plot on the left displays the UMAP scatterplot for the front view of resting sheep, while the plot on the right shows the side view. In each plot, data points represent individual instances of sheep resting postures. The clustering of points in the UMAP space signifies patterns of similarity within the resting postures. A tighter clustering indicates a more consistent and homogeneous behavior pattern among the animals within that group. From these plots, we observe that sheep demonstrate a propensity for maintaining a consistent resting posture when alone, which is reflected in the tightly clustered gray points. This homogeneity in behavior could be attributed to a lack of external stimuli, leading to a uniformity in resting positions. In contrast, the brown clusters corresponding to sheep resting in a herd show a more scattered distribution, suggesting a diversity in resting postures. This variability could arise from social dynamics within the herd or because of spatial constraints. The scattered clusters imply that sheep in a herd may change their resting positions more frequently or adopt different postures compared to when they are alone. The UMAP plots provide a compelling visual summary of the resting behaviors, highlighting the differences between solitary and social resting states in sheep. This insight is crucial for understanding sheep behavior in controlled environments and has implications for animal welfare and the management of livestock. These UMAP visualizations eloquently encapsulate the dichotomy between solitary and social resting states among sheep. The implications of this data are profound, aligning with the grazing behavior findings and reinforcing the notion that sheep display more uniform behaviors when solitary. Such insights are pivotal for the realm of animal husbandry, particularly in the context of enhancing livestock management and welfare practices.

5. Conclusion

Our study presents a robust multimodal AI-based vision framework that integrates cutting-edge models for comprehensive behavioral analytics applicable to precision livestock farming. Through the analysis of a sample sheep dataset, we have showcased the utility of our framework in extracting detailed insights into various sheep activities such as running, grazing, and resting. The application of our framework has revealed fascinating observations that are of significant value to the agricultural and farming communities. For instance, the solitary grazing patterns of sheep in a controlled environment without apparent dangers suggest potential adjustments in farm management practices. The insights into their resting patterns further emphasize the impact of social dynamics on animal behavior, which could guide enhancements in livestock care and facility design. Similarly, our insights from sheep running videos validate the finding that repetitive movements result in faster speed [1]. However, the current dataset did not encompass certain activities, such as milking, and these findings are based on brief observational videos, limiting the scope of our analyses to report their health. Additionally, while we focused on sheep, extending this framework to other species (e.g. cows) could provide a broader understanding of livestock behaviors and welfare across different farming systems.

In future, our goal is to apply this framework to diverse species and environmental conditions. This will enable us to not only continue our exploration of animal activities but also investigate aspects such as dietary habits in relation to the nutritional content of forage. By correlating behavior with nutritional intake, future work could offer guidance on pasture management to optimize feeding efficiency and herd health. We believe AnimalFormer has the potential to revolutionise livestock monitoring by providing non-invasive, in-depth behavioral analysis without any additional hardware (e.g. RFID-based animal tagging). Such a multimodal vision framework is indispensable for advancing precision farming operations, promoting animal welfare, and improving the productivity and sustainability of the agricultural sector.

References

- Gaia Ariani, Young-Hoon Kwon, and Jörn Diedrichsen. Repetita iuvant: repetition facilitates online planning of sequential movements. *Journal of Neurophysiology*, 123(5): 1727–1738, 2020. 7, 8
- Shaun Barney, Satnam Dlay, Andrew Crowe, Ilias Kyriazakis, and Matthew Leach. Deep learning pose estimation for multi-cattle lameness detection. *Scientific Reports*, 13(1): 4499, 2023. 3
- [3] Alain Boissy and Bertrand Dumont. Interactions between social and feeding motivations on the grazing behaviour of herbivores: sheep more easily split into subgroups with familiar peers. *Applied Animal Behaviour Science*, 79(3):233–245, 2002. 8
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. arXiv preprint arXiv:2104.14294, 2021. 4
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. arXiv preprint arXiv:2002. 4
- [6] Richard Cooper and Françoise Wemelsfelder. Qualitative behaviour assessment as an indicator of animal emotional welfare in farm assurance. *Livestock*, 25(4):180, 2020. 1
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 3
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*, 2020. 2
- [9] CM Dwyer. How has the risk of predation shaped the behavioural responses of sheep to fear and distress? *Animal Welfare*, 13(3):269–281, 2004. 7
- [10] Cheng Fang, Tiemin Zhang, Haikun Zheng, Junduan Huang, and Kaixuan Cuan. Pose estimation and behavior classification of broiler chickens based on deep neural networks. *Computers and Electronics in Agriculture*, 180:105863, 2021. 1
- [11] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning, 2020. 4
- [12] Yaning Han, Ke Chen, Yunke Wang, Wenhao Liu, Zhouwei Wang, Xiaojing Wang, Chuanliang Han, Jiahui Liao, Kang Huang, Shengyuan Cai, et al. Multi-animal 3d social pose estimation, identification and behaviour embedding with a few-shot learning framework. *Nature Machine Intelligence*, pages 1–14, 2024. 3
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn, 2018. 1
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual rep-

resentation learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 4

- [15] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021. 3
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. 4
- [17] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality, 2023. 1, 3, 5
- [18] Nathan A. Kelly, Bilal M. Khan, Muhammad Y. Ayub, Abir J. Hussain, Khalil Dajani, Yunfei Hou, and Wasiq Khan. Video dataset of sheep activity for animal behavioral analysis via deep learning. *Data in Brief*, 2024(110027), 2024. Under a Creative Commons license. 1, 3
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023. 5
- [20] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. Grounding dino: Marrying dino with grounded pre-training for open-set object detection, 2023. 1, 3, 4
- [21] Wolfgang Lutz and Samir K C. Dimensions of global population projections: what do we know about future population trends and structures? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1554):2779– 2791, 2010. 1
- [22] James B MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, 1(14):281–297, 1967. 5
- [23] Markus Marks, Qiuhan Jin, Oliver Sturman, Lukas von Ziegler, Sepp Kollmorgen, Wolfger von der Behrens, Valerio Mante, Johannes Bohacek, and Mehmet Fatih Yanik. Deeplearning-based identification, tracking, pose estimation and behaviour classification of interacting primates and mice in complex environments. *Nature machine intelligence*, 4(4): 331–340, 2022. 1, 3
- [24] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426, 2018. 5
- [25] Jeffrey S Mogil, Daniel S Pang, Giorgia G S Dutra, and Christine T Chambers. The development and use of facial grimace scales for pain measurement in animals. *Neuroscience & Biobehavioral Reviews*, 116:480–493, 2020. 1
- [26] Daniel Mota-Rojas, Donald M Broom, Agustín Orihuela, Antonio Velarde, Fabio Napolitano, and María Alonso-Spilsbury. Effects of human-animal relationship on animal productivity and welfare. *Journal of Animal Behaviour and Biometeorology*, 8:196–205, 2020. 1
- [27] Suresh Neethirajan. Affective state recognition in livestock—artificial intelligence approaches. *Animals*, 12:759, 2022. 1

- [28] Suresh Neethirajan, Inonge Reimert, and Bas Kemp. Measuring farm animal emotions—sensor-based approaches. *Sensors*, 21(2):553, 2021. 1
- [29] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In 2014 IEEE Conference on Computer Vision and Pattern Recognition. IEEE, 2014. 1
- [30] Alice V Viscardi, Matthew Hunniford, Patrick Lawlis, Matthew Leach, and Patricia V Turner. Development of a piglet grimace scale to evaluate piglet pain using facial expressions following castration and tail docking: A pilot study. *Frontiers in Veterinary Science*, 4:51, 2017. 1
- [31] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In Advances in Neural Information Processing Systems, 2022. 1, 3
- [32] Chen Yang, Jeremy Forest, Matthew Einhorn, and Thomas A. Cleland. Automated behavioral analysis using instance segmentation, 2023. 1, 3
- [33] Yuxiang Yang, Junjie Yang, Yufei Xu, Jing Zhang, Long Lan, and Dacheng Tao. Apt-36k: A large-scale benchmark for animal pose estimation and tracking, 2022. 1
- [34] Shaokai Ye, Anastasiia Filippova, Jessy Lauer, Steffen Schneider, Maxime Vidal, Tian Qiu, Alexander Mathis, and Mackenzie Weygandt Mathis. Superanimal pretrained pose estimation models for behavioral analysis, 2023. 1, 3
- [35] Hang Yu, Yufei Xu, Jing Zhang, Wei Zhao, Ziyu Guan, and Dacheng Tao. Ap-10k: A benchmark for animal pose estimation in the wild, 2021. 1
- [36] Kai Zhang, Shuqing Han, Jianzhai Wu, Guodong Cheng, Yali Wang, Saisai Wu, and Jifang Liu. Early lameness detection in dairy cattle based on wearable gait analysis using semi-supervised lstm-autoencoder. *Computers and Electronics in Agriculture*, 213:108252, 2023. 3