

Reviving the Context: Camera Trap Species Classification as Link Prediction on Multimodal Knowledge Graphs

Vardaan Pahuja
pahuja.9@osu.edu
The Ohio State University
Columbus, Ohio, United States

Weidi Luo
luo.1455@osu.edu
The Ohio State University
Columbus, Ohio, United States

Yu Gu
gu.826@osu.edu
The Ohio State University
Columbus, Ohio, United States

Cheng-Hao Tu
tu.343@osu.edu
The Ohio State University
Columbus, Ohio, United States

Hong-You Chen
chen.9301@osu.edu
The Ohio State University
Columbus, Ohio, United States

Tanya Berger-Wolf
berger-wolf.1@osu.edu
The Ohio State University
Columbus, Ohio, United States

Charles Stewart
stewart@rpi.edu
Rensselaer Polytechnic Institute
Troy, New York, United States

Song Gao
song.gao@wisc.edu
University of Wisconsin-Madison
Madison, Wisconsin, United States

Wei-Lun Chao
chao.209@osu.edu
The Ohio State University
Columbus, Ohio, United States

Yu Su
su.809@osu.edu
The Ohio State University
Columbus, Ohio, United States

Abstract

Camera traps are important tools in animal ecology for biodiversity monitoring and conservation. However, their practical application is limited by issues such as poor generalization to new and unseen locations. Images are typically associated with diverse forms of context, which may exist in different modalities. In this work, we exploit the structured context linked to camera trap images to boost out-of-distribution generalization for species classification tasks in camera traps. For instance, a picture of a wild animal could be linked to details about the time and place it was captured, as well as structured biological knowledge about the animal species. While often overlooked by existing studies, incorporating such context offers several potential benefits for better image understanding, such as addressing data scarcity and enhancing generalization. However, effectively incorporating such heterogeneous context into the visual domain is a challenging problem. To address this, we propose a novel framework that transforms species classification as link prediction in a multimodal knowledge graph (KG). This framework enables the seamless integration of diverse multimodal contexts for visual recognition. We apply this framework for out-of-distribution species classification on the iWildCam2020-WILDS and Snapshot Mountain Zebra datasets and achieve competitive performance with

state-of-the-art approaches. Furthermore, our framework enhances sample efficiency for recognizing under-represented species.¹

CCS Concepts

• **Computing methodologies** → **Semantic networks; Object recognition.**

Keywords

Camera Traps, Species Classification, Multimodal Knowledge Graph, KG Link Prediction

ACM Reference Format:

Vardaan Pahuja, Weidi Luo, Yu Gu, Cheng-Hao Tu, Hong-You Chen, Tanya Berger-Wolf, Charles Stewart, Song Gao, Wei-Lun Chao, and Yu Su. 2024. Reviving the Context: Camera Trap Species Classification as Link Prediction on Multimodal Knowledge Graphs. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Human activities are increasingly endangering wildlife species, resulting in a significant global decline in animal populations [2, 19, 37]. Therefore, accurately identifying and tracking wildlife species is vital for preserving ecological biodiversity. Camera traps, digital cameras activated by motion or infrared in natural habitats, have become ecologists' preferred data collection tool [23, 44, 67]. However, manually sifting through the numerous images they capture is a time-consuming and arduous task for experts. This has led to the increased use of computer vision techniques for species recognition [1, 13, 30, 52, 56, 68]. Yet, a challenge has arisen: many of these

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

¹Our code is available at <https://github.com/OSU-NLP-Group/COSMO>

models overfit to the backgrounds of their training images, diminishing their effectiveness on images from new locations [9, 39, 55]. This underscores the need for *more adaptable species classification models that perform well across diverse contexts*.

Building on this, cognitive science research has demonstrated the profound influence of **contextual** information on human perception and visual recognition processes [4, 5, 45]. Particularly in wildlife monitoring, camera trap images are replete with crucial contextual data, such as where (*i.e.*, camera location coordinates) and when (*i.e.*, timestamps) a photo is taken. Furthermore, the structured knowledge of biology taxonomy (*e.g.*, Open Tree Taxonomy [46]) can also provide valuable context for understanding the species in camera trap images. Such context provides important knowledge that can boost the recognition of visual concepts. For instance, the knowledge that a certain feline image was taken from a camera trap in Africa significantly reduces the likelihood of it representing a tiger. In addition, more robust associations might be learned with the aid of contextual information because the context provides invariable knowledge that is unbiased towards variations in the illuminations or angles of an image. This may help to compensate for domain shifts in species images resulting from such variations and potentially lead to better out-of-distribution (OOD) generalizability [6, 21]. Consequently, the incorporation of contextual information in species identification presents a significant problem worthy of investigation.

Nevertheless, contextual information has been under-exploited in the literature of image classification; standard image classification models [24, 59] often disregard the contextual information tied to images. This is partly due to the heterogeneous nature of the context, which makes it challenging to incorporate contextual information in image classification using a unified learning framework. Contextual information in different modalities (*e.g.*, numerical values, textual descriptions, or structured taxonomies) is usually represented separately from the image in *distinct feature spaces*. The question of effectively combining features from these different spaces within a unified learning framework remains unanswered. Existing research typically treats all the features as additional input to the classifier via feature vector concatenation [6, 21, 32] or utilizes fusion to obtain aggregate representations [16, 18]. Despite their simplicity, such approaches are incapable of capturing complex structural and semantic relationships between images and various contextual information. Additionally, these approaches assume a uniform availability of contextual information across all images, which is often unrealistic in real-world scenarios. As a result, their flexibility is limited, especially when considering situations where certain images may lack some contextual details, such as coordinates or timestamps, like in camera trap photos.

Towards this end, we propose a new learning framework, COSMO (Classification Of Species using Multimodal cOntext), where we first organize all species images and contextual information as a *multimodal knowledge graph* (KG) and then reformulate species classification as the standard link prediction task on the KG. Specifically, we consider species images, their corresponding labels (which are available in the training data), and their associated attributes provided in the context as entities within our KG (see Figure 1 for an example). We represent the relationships between these entities

as edges in our KG (see a more concrete description of our KG construction in Section 3.2). Our KG is multimodal because its entities belong to different modalities. In this context, species classification can be framed as a link prediction task, where the objective is to predict the presence of an edge between an image and its corresponding species label within the KG. This learning framework enables a unified way to incorporate heterogeneous contextual information for species classification. Each form of multimodal information is treated as a type of entity, a first-class citizen of the multimodal KG with its representation computed using a modality-specific encoder. The learning process enables the interaction of different modalities in a *joint feature space* for robust representation learning. In addition, COSMO demonstrates greater flexibility by not assuming uniform availability of all contextual information, unlike previous methods.

We employ the widely used DistMult [70] model as our backbone model for link prediction to instantiate the COSMO framework. To assess the performance of COSMO, particularly in terms of out-of-distribution generalization, we conduct experiments on the iWildCam2020-WILDS benchmark [30] and Snapshot Mountain Zebra [48], which are standard datasets for species classification in camera trap photos. They contain naturally occurring wildlife photos associated with metadata. Factors like variation in illumination, camera pose, and motion blur pose challenges for robustness and generalization, making these benchmarks an ideal testbed for assessing our framework's effectiveness. We show that COSMO offers a unified framework to incorporate heterogeneous context leading to improved species classification performance over existing out-of-distribution generalization approaches.

The main contribution of this work is three-fold:

- We propose a novel framework, COSMO, that reformulates species classification as link prediction in a multimodal knowledge graph, which provides a unified way to incorporate heterogeneous forms of contextual information associated with images for visual recognition.
- We instantiate this framework for species classification of wildlife images, including the construction of a novel multimodal KG for this problem that integrates spatiotemporal information and structured biology knowledge.
- Evaluation on the standard iWildCam2020-WILDS and Snapshot Mountain Zebra datasets demonstrate that COSMO achieves competitive performance compared with standard species classification methods, especially in improving robustness and OOD generalization.

2 Related Work

Species Recognition in Camera Traps. Deep neural networks such as CNNs have been successfully deployed for large-scale recognition of camera trap images [43, 64, 68]. This has paved the way for significant savings in logistics costs for biodiversity conservation. However, training such models often requires enormous amounts of data to perform well. Sadegh Norouzzadeh et al. and Bothmann et al. propose active learning approaches to mitigate the sample inefficiency of training species classification models in such systems. Another challenge arises from the tendency of




Figure 1: Overview of our framework COSMO. Left: Our multimodal knowledge graph for camera traps and wildlife. Photos from camera traps are jointly represented in the KG with contextual information such as time, location, and structured biology taxonomy. The taxonomy is obtained from Open Tree Taxonomy (OTT) [46] or iNaturalist [25]. **Right:** In our formulation of species classification as link prediction, the plausibility score $\psi(s, r, o)$ of each (subject, relation, object) triplet is computed using a KGE model (e.g., DistMult), where the subject, relation, and object are all first embedded into a vector space. Specifically, for our multimodal KG, we represent visual entities using a ResNet-50 pre-trained on ImageNet and represent numerical entities using an MLP. For categorical entities and relations, we directly represent them with embedding lookups.

these models to overfit to the backgrounds present in the training images [9, 39], which limits their deployment to new camera trap locations [55, 64]. Improving robustness to new locations is a significant research challenge [8, 63] leading to the curation of datasets like iWildCam2020-WILDS [30] to test OOD generalization for such systems. Domain adaptation approaches in the literature seek to mitigate this issue by distributionally robust optimization [26, 50] or learning domain invariant features [61]. In contrast, this work helps improve the robustness to new camera trap locations by utilizing a multimodal KG of heterogeneous contexts.

Image Classification with Auxiliary Information. Despite the ubiquity of contextual metadata, the potential of leveraging them for image classification has been largely under-explored. Previous studies have primarily treated metadata as additional input features for classifiers [6, 21, 32], representing a shallow use that fails to capture the intricate relationships between metadata and images. Some works have attempted to model pairwise dependencies between images using heuristics based on metadata, such as shared tags on social media [35, 38] or aggregating information from neighborhood images with similar metadata [28] while disregarding more complex relationships among images, metadata, and labels. Metaformer [18] feeds a sequence of image patches and metadata to a Transformer model for their fusion. Additionally, these methods assume a uniform availability of metadata for all images, which is often not the case in reality due to data scarcity. For

instance, the camera trap location coordinates may not be available in some cases due to privacy and security reasons. In our work, we do not assume such uniform availability and build the multimodal KG using available metadata.

Apart from the metadata, external sources of knowledge are also used in image classification. For instance, Jayatilaka et al. embed each class as a vector based on a hierarchy derived from WordNet [40]. Alsallakh et al. develop a class hierarchy-aware CNN for image classification on ImageNet. Similarly, Bertinetto et al. and Zhang et al. design hierarchy-aware objectives to incorporate taxonomy in image representations. BioCLIP [60] verbalizes the taxonomic hierarchy to train a CLIP-style foundational model for species classification across plants, animals, and fungi. Marino et al. represent images as local subgraphs of Visual Genome [31]. In contrast, COSMO constructs a global KG with both metadata and external knowledge, e.g., taxonomy information from Open Tree Taxonomy, and approaches image classification as link prediction within the KG. Our novel formulation is flexible in handling data scarcity of metadata and enables reasoning over diverse relationships present in the KG.

KG Link Prediction. Most real-world KGs are incomplete. The task of link prediction or knowledge graph completion (KGC) tries to infer missing links given the observed ones. Early approaches for link prediction range from translation-based models [12, 34] and semantic matching models [42, 70] to the ones that leverage neural

networks like feedforward neural networks [20], CNNs [17, 41], and Transformer-based models [15, 53, 71]. These methods use a parameterized scoring function based on learned entity and relation embeddings to calculate the plausibility of a particular triplet. However, it could be challenging to fully encode the rich semantic information of KGs into such shallow embeddings. To mitigate this, Schlichtkrull et al., Vashishth et al., Yu et al., Pahuja et al. use graph neural networks (GNNs) to encode the rich neighborhood context of entities for link prediction. In our framework, we employ a global multimodal KG, which consists of biological taxonomy and metadata, as the context to enhance OOD generalization.

Multimodal KG Reasoning. Multimodal KGs extend traditional KGs by including entities of different modalities such as categorical data, images, numerical data, etc. KBLRN [22] is a pioneering work in multimodal KG reasoning that uses extra information in the form of relational and numerical features for multimodal KG reasoning. Similarly, IKRL [57] proposes a fusion of linguistic and visual information with structured information for link prediction. MKBE [49] constructs a multimodal KG using numerical, image, and textual information, treating them as entities instead of auxiliary features, for the link prediction task. MR-GCN [69] further extends it by including support for more modalities, *e.g.*, numerical, temporal, textual, visual, and spatial predicate links in the multimodal KG. To provide a more expressive way for interaction between different modalities, IMF [33] uses bilinear pooling to fuse multiple modality features and trains it using contrastive learning on the contextual entity representations. Our work leverages link prediction in a multimodal KG to enable out-of-distribution generalization for species classification in camera traps.

3 Methodology

3.1 Preliminaries

Multimodal KG. Given a set of KG entities with categorical values $\mathcal{E}_{\mathcal{KG}}$, multimodal entities $\mathcal{E}_{\mathcal{MM}}$, and a set of relations $\mathcal{R}$, a multimodal KG can be defined as a collection of facts $\mathcal{F} \subseteq (\mathcal{E}_{\mathcal{KG}} \cup \mathcal{E}_{\mathcal{MM}}) \times \mathcal{R} \times (\mathcal{E}_{\mathcal{KG}} \cup \mathcal{E}_{\mathcal{MM}})$ where for each fact $f = (h, r, t)$, $h, t \in (\mathcal{E}_{\mathcal{KG}} \cup \mathcal{E}_{\mathcal{MM}}), r \in \mathcal{R}$.

KG Link Prediction. The task of link prediction is to infer missing facts based on known facts in a KG. Given a link prediction query $(h, r, ?)$ or $(?, r, t)$, the model ranks the target entity among the set of candidate entities.

Problem Setup. The task entails species recognition for camera trap images amidst distribution shifts. The training and test sets comprise images obtained from disjoint camera traps, enabling the evaluation of out-of-domain (OOD) generalization. During training, we use the multimodal KG to train our model, while we use just the image to make predictions for inference. The goal is to learn visual representations robust to distribution shifts by leveraging the rich structural and semantic information provided by the multimodal knowledge graph.

3.2 Building the Multimodal KG

The multimodal KG comprises entities from different modalities interconnected by heterogeneous relationships. The base KG consists of camera trap images linked with their species labels from

the training set ($\langle \text{image} \rangle$, instance of, $\langle \text{species label} \rangle$). Next, we progressively augment the KG with links connecting the existing entities to contextual information. In this work, we utilize the following attributes to provide context for species classification:

- **Taxonomy:** The taxonomy forms the core of the multimodal knowledge graph, connecting distinct species to higher-order taxa. For iWildCam2020-WILDS, we obtain the phylogenetic taxonomy corresponding to the species of interest from Open Tree Taxonomy (OTT) [46] and manually link it to the species in the dataset. For the Snapshot Mountain Zebra dataset, we utilize the iNaturalist taxonomy [25] mapping provided by www.lila.science.
- **Location:** The camera trap images are associated with the GPS coordinates of their source cameras. For the iWildCam2020-WILDS dataset, this metadata is available for a portion of the images (67%) and is obfuscated within 1 km. for privacy reasons. Animals demonstrate a preference for particular habitats; thus, the location context attribute is useful for species recognition.
- **Time:** The timestamp attribute indicates the precise moment when the image was captured. This timestamp information proves valuable in species recognition since specific animals exhibit activity patterns tied to particular times of the day, such as feeding, hunting, or defending their territory. In our multimodal knowledge graph, we utilize the timestamp information at an hourly granularity.

Figure 1 presents a schematic representation of various contexts in a multimodal KG. For location, time, and taxonomy attributes, the corresponding RDF triplets can be represented as $\langle \text{image} \rangle$, location, $\langle \text{GPS co-ordinate} \rangle$, $\langle \text{image} \rangle$, time, $\langle \text{timestamp} \rangle$, and $\langle \text{taxon}_1 \rangle$, parent, $\langle \text{taxon}_2 \rangle$, respectively.

3.3 Model Architecture

We use DistMult [70], a strong baseline on KGE benchmarks, as our backbone KG embedding model.² Note that COSMO is a general framework that can leverage a variety of KG embedding models proposed in the literature. DistMult minimizes a bilinear scoring function between the entity embeddings of the subject and object entities. For a given triplet (h, r, t) , the scoring function of DistMult is defined as:

$$\psi(h, r, t) = \mathbf{h}^T \mathbf{W}_r \mathbf{t} = \sum_{i=1}^d \mathbf{h}_i \cdot \text{diag}(\mathbf{W}_r)_i \cdot \mathbf{t}_i \quad (1)$$

Here, $\mathbf{h}$ and $\mathbf{t}$ denote the vector representations of the head entity and tail entity, respectively. The relation representation is parameterized by $\mathbf{W}_r \in \mathbb{R}^{d \times d}$, a diagonal matrix.

3.3.1 Multi-modality Encoders. We use an ImageNet pre-trained ResNet-50 [24] as the image encoder. The base feature of each location is represented as a 2D vector [latitude, longitude]. Following prior work [49], we use an MLP to project the 2D location feature to a higher dimensional space. Similarly, for temporal context, we use an MLP to project the integer value of the hour

²Recent work [51] showed that simple baselines like DistMult outperform more sophisticated neural network baselines when trained properly.

timestamp to the higher dimensional embedding space. For categorical entities such as species labels and taxa, we learn dense embeddings as representations.

3.3.2 Training. We train the model using an optimization strategy based on the modality of the tail entity. For categorical attributes, we formulate it as a multi-class classification problem and use standard cross-entropy loss to train the model. For instance, in case of a given image-species label ground truth triplet $(I, \text{instance of}, s)$, the loss is defined as:

$$\mathcal{L}(I, \text{instance of}, s) = -\log \frac{\exp(\psi(I, \text{instance of}, s))}{\sum_{s' \in S} \exp(\psi(I, \text{instance of}, s'))}, \quad (2)$$

where S denotes the set of all species labels and $\psi(h, r, t)$ denotes the plausibility score of KG edge (h, r, t) .

For numerical attributes such as location and time, we formulate it as a multi-class multi-label classification problem and use a binary cross-entropy loss to optimize the parameters. This choice is motivated by the fact that images can be associated with a range of GPS coordinates and timestamps, e.g., most animals are active multiple times during the day. The label space comprises all entities of ground truth modality. For instance, in the case of a given time modality ground truth triplet (I, time, t) , the loss is defined as:

$$\mathcal{L}(I, \text{time}, t) = -\sum_{t'} l_{t'}^{I, \text{time}} \cdot \log(\sigma(\psi(I, \text{time}, t'))) + (1 - l_{t'}^{I, \text{time}}) \cdot (1 - \log(\sigma(\psi(I, \text{time}, t')))), \quad (3)$$

where $l_{t'}^{I, \text{time}}$ is a binary label that indicates whether the triplet (I, time, t') exists in the set of observed triplets and $\sigma(\cdot)$ is the sigmoid activation function. We train the model by sequentially minimizing the objective on each type of context triplet. Figure 1 illustrates the overall model architecture.

4 Experimental Setup

4.1 Datasets

We test our approach on the iWildCam2020-WILDS dataset [30], a variant of the iWildCam 2020 dataset [7] and Snapshot Mountain Zebra [48]. iWildCam2020-WILDS is a benchmark dataset designed to test OOD generalization for the task of species classification. The label space consists of 182 species. Each domain corresponds to a different location of the camera trap. The training and test images belong to disjoint sets of locations in the OOD setting.

Snapshot Mountain Zebra comprises camera trap images taken at the Mountain Zebra National Park in South Africa as a part of the Snapshot Safari project [48]. The label space consists of 53 species, mostly annotated at the species level. Prominent animal species include Cape Mountain zebra, kudu, and springbok. The location coordinates are not available for this dataset due to privacy and security reasons. We manually split the images to have disjoint camera traps in each split due to the absence of a standard split. These datasets pose a significant challenge for species recognition due to factors like inadequate illumination, motion blur, occlusion, temporal variations, and diverse weather conditions, effectively reflecting the complexities of real-life camera trap usage. Dataset statistics are shown in Table 2.

4.2 Baselines

We use the COSMO with no context that uses just the species label edges as our baseline. In addition, we compare with the following baseline algorithms for OOD generalization: Empirical Risk Minimization (ERM) [30], which trains the model to minimize average training loss, CORAL [61], a method for unsupervised domain adaptation that learns domain invariant features, Group DRO [26], an algorithm that uses distributionally robust optimization to perform well on subpopulation shifts, Fish [58] that attempts domain adaptation using gradient matching, and ABSGD [50], an optimization method for addressing data imbalance. As an alternative way of incorporating contextual information, we implement MLP-concat, a baseline which utilizes the location and temporal features at both training and inference time. It uses vanilla concatenation to fuse visual and spatiotemporal representations which are then fed into an MLP. The missing features are substituted by a mean value computed over the training dataset. All models use a pre-trained ResNet-50 as image encoder. We evaluate the models using overall accuracy as the metric.

4.3 Implementation Details

We implement our models in PyTorch. The hidden dimension of the multimodal KG embedding model is set to 512, with a batch size of 16. The images are resized to 448×448 before input to the image encoder. For the location and time attributes, we use a 3-layer MLP that projects the feature input dimension to the embedding dimension and uses PReLU as the activation function. We use Adam [29] optimizer with a learning rate of $3e-5$ and $1e-3$ for the image encoder and the rest of the parameters, respectively. In our experiments, the models on iWildCam2020-WILDS and Snapshot Mountain Zebra were trained for 12 and 15 epochs, respectively. We use early stopping based on validation accuracy to prevent overfitting. The early stopping patience parameter is set to 5 epochs. All results are reported with averages across three random seeds.

5 Results

In this section, we attempt to answer the following questions:

- Q1. Does the use of contextual information contribute toward better performance? (Section 5.1)
- Q2. How does COSMO's performance compare to the existing state-of-the-art? (Section 5.2)
- Q3. Does the taxonomy-aware COSMO model result in more semantically plausible predictions? (Section 5.4.1)
- Q4. How does COSMO's performance compare to baselines for under-represented species? (Section 5.4.3)

5.1 Performance Comparison with Addition of Multimodal Context

We add taxonomy, location, and temporal context information to the base KG and observe the impact on the species classification performance. Table 1 shows the results for the iWildCam2020-WILDS dataset. We make the following observations from these results:

Firstly, the addition of one or more contexts results in a performance gain over the no-context baseline in the vast majority of cases. For instance, in the case of COSMO with taxonomy, we obtain

Table 1: Species classification results on iWildCam2020-WILDS (OOD) dataset. The first baseline in the second section shows the no-context baseline that uses only image-species labels as KG edges. All models use a pre-trained ResNet-50 as image encoder. Parentheses show standard deviation across 3 random seeds. We highlight the best result in bold and the second best with underline. We mark the improvements over COSMO (no-context) in **green. Missing values are denoted by –.**

	Model	Multi-modality			Val. Acc. (%)	Test Acc. (%)
		Taxonomy	Location	Time		
	Empirical Risk Minimization (ERM) [30]				62.7 (± 2.4)	71.6 (± 2.5)
	CORAL [61]				60.3 (± 2.8)	73.3 (± 4.3)
	Group DRO [26]		–		60.0 (± 0.7)	72.7 (± 2.0)
	Fish [58]				58.0 (± 0.2)	63.2 (± 0.7)
	ABSGD [50]				–	72.7 (± 1.8)
	MLP-concat		✓	✓	27.3 (± 0.8)	39.6 (± 1.0)
	COSMO (no-context)		–		63.2 (± 0.4)	68.8 (± 2.1)
Single context	COSMO	✓			62.8 (± 2.2) (–0.4)	72.4 (± 2.5) (+3.6)
			✓		64.4 (± 1.0) (+1.2)	74.5 (± 3.6) (+5.7)
				✓	64.7 (± 0.4) (+1.5)	71.1 (± 3.1) (+2.3)
Multiple contexts	COSMO	✓	✓		65.4 (± 0.4) (+2.2)	70.4 (± 2.1) (+1.6)
		✓		✓	<u>64.9</u> (± 1.6) (+1.7)	73.7 (± 3.8) (+4.9)
			✓	✓	63.0 (± 2.1) (–0.2)	<u>74.2</u> (± 2.2) (+5.4)
		✓	✓	✓	65.0 (± 1.6) (+1.8)	71.5 (± 2.8) (+2.7)

Table 2: Dataset Statistics.

Dataset	Split	# Images	# Camera traps
iWildCam2020-WILDS	Train	129,809	243
	Val.	14,961	32
	Test	42,791	48
Snapshot Mountain Zebra	Train	39,820	13
	Val.	14,754	3
	Test	18,417	3

a 3.6% improvement over the no-context baseline in terms of test accuracy. Incorporating location context produces a notable 5.7% enhancement in test set accuracy, underlining the significance of auxiliary information for improved out-of-domain generalization. We further analyze the role of location in predicting the species distribution in Section 5.4.2. Additionally, utilizing the time attribute yields a substantial improvement over the no-context baseline, resulting in a 2.3% performance gain.

Secondly, we observe that the use of multiple contexts results in a performance boost in a majority of cases. For instance, the addition of location and time attributes improves over the taxonomy baseline by a margin of 2.6% and 2.1% respectively in terms of the validation set accuracy. Similarly, the taxonomy with time baseline obtains an improvement of 1.3% and 2.6% over the taxonomy and time baselines, respectively in terms of test accuracy.

Table 3 shows the results for the Snapshot Mountain Zebra dataset. Incorporating taxonomy and time contexts results in a

Table 3: Species classification results on Snapshot Mountain Zebra dataset. We obtain the results for OOD baselines by training them on this dataset using publicly available code. We mark the improvements over COSMO (no-context) in **green.**

Model	Multi-modality		Test Acc. (%)
	Taxonomy	Time	
ERM [30]			96.2 (± 0.6)
CORAL [61]			96.6 (± 1.2)
Group DRO [26]	–		93.4 (± 2.1)
ABSGD [50]			93.4 (± 2.0)
MLP-concat		✓	94.7 (± 0.0)
COSMO (no-context)	–		92.9 (± 2.5)
COSMO	✓		93.9 (± 2.8) (+1.0)
		✓	95.3 (± 3.1) (+2.4)
	✓	✓	96.8 (± 0.4) (+3.9)

performance boost of 1% and 2.4% respectively, over the no-context baseline. Furthermore, their combined use yields a noteworthy 3.9% gain in test accuracy.

5.2 Comparison with OOD Generalization Approaches

We compare the performance of the COSMO with methods specifically designed for out-of-domain generalization. Notably, our best-performing model, which uses location as context, achieves state-of-the-art performance in terms of OOD test accuracy, outperforming the existing SOTA model (CORAL) by 1.2% on the iWildCam2020-WILDS dataset. Likewise, COSMO with taxonomy and time contexts outperforms existing approaches on the Snapshot Mountain Zebra dataset. This demonstrates the effectiveness of leveraging diverse multimodal contexts for achieving more robust OOD generalization, even in the absence of sophisticated objectives aimed at improving domain generalization, *e.g.*, CORAL, Group DRO, ABSGD, and Fish. The MLP-concat baseline overfits the training camera trap locations on the iWildCam2020-WILDS dataset, resulting in suboptimal performance. COSMO consistently outperforms the MLP-concat baseline by a significant margin across both datasets.

Table 4: Performance comparison among different KGE models (iWildCam2020-WILDS). DistMult outperforms ConvE in a majority of cases.

Model Configuration	KGE	Val. Acc.
No-context	DistMult	63.2 (0.4)
	ConvE	62.2 (0.4)
Taxonomy Only	DistMult	62.8 (2.2)
	ConvE	64.7 (0.6)
Location Only	DistMult	64.4 (1.0)
	ConvE	63.8 (1.8)
Time Only	DistMult	64.7 (0.4)
	ConvE	59.6 (2.8)
Taxonomy, Location, Time	DistMult	65.0 (1.6)
	ConvE	55.8 (3.5)

Table 5: Quantitative evaluation of COSMO errors with and without taxonomy using a hierarchical distance metric (iWildCam2020-WILDS). The taxonomy-aware model achieves better performance in terms of Avg. LCA height.

Model	Avg. LCA height
COSMO (no-context)	6.72
COSMO (w/ taxonomy)	6.39

5.3 Comparison with Alternative KGE Backbones

In our preliminary experiments, we explored the use of ConvE [17], a strong neural network baseline, as an alternative to DistMult for the KGE backbone model for the iWildCam2020-WILDS dataset (Table 4). Sun et al. [62] show that ConvE outperforms more recent neural network KGE models when evaluated properly. We observe that DistMult outperforms ConvE in a majority of cases, particularly

when all incorporating all context types simultaneously. Furthermore, DistMult offers the advantage of being more computationally efficient than neural network based KG embedding approaches.

5.4 Fine-grained Analyses

5.4.1 Error Analysis for the Taxonomy-aware Model. To analyze the predictions of our model with and without taxonomy information, we employ a metric that takes into account the hierarchical structure of the labels. Conventional measures like top-1 accuracy treat all errors equally, disregarding the semantic relationships among labels. Hence, we use the Least Common Ancestor (LCA) [10] for the misclassified examples as the metric for this analysis (Table 5). A lower LCA value indicates that the errors made by the taxonomy-aware model are more semantically related to the true label compared to the baseline.

We compare the predictions of COSMO which uses taxonomy to the no-context baseline (Figure 2). Notably, the inclusion of taxonomy information assists the model in avoiding implausible predictions. For instance, consider the case of the animal ocelot (*Leopardus pardalis*), which belongs to the cat family (*feliformia*). The use of taxonomy information prevents the misprediction of this animal as a gray fox, which belongs to the dog family (*caniformia*). Similarly, in the second example, the baseline model incorrectly predicts the given image as Central American agouti, a mammal, instead of ocellated turkey, a bird.

5.4.2 Correlation Analysis for Location and Time Attributes. We examined the relationship between species distribution and numerical attributes, such as location and time, to gain insights into how these contexts contribute to the task. The location coordinates can be grouped into six clusters. A visualization of the location clusters is shown in Figure 4. For each pair of cluster centroids, we compute the Bhattacharyya distance [11], a measure of similarity between probability distributions, between the training and validation set species distributions (Figure 3a). Similarly, we plot the distance between species distributions corresponding to each hour of the day (Figure 3b). We observe that the similarity (corresponds to lower distance) peaks along the diagonal for the location attribute, as well as for the day/night categorization of the time attribute. This suggests these metadata give a prior for species class distribution.

Table 6: Performance comparison on under-represented species classification (OOD test set). The best performing COSMO model improves over the ERM baseline by a significant margin.

Model	Accuracy
ERM	16.3
COSMO	19.0 (+2.7)

³The null value in row 4 is due to the absence of species overlap with respective validation clusters. The null value in columns 3 and 4 indicates the absence of these clusters in the validation set.




Figure 2: Comparison of COSMO model with and without taxonomy edges (iWildCam2020-WILDS validation set). The use of taxonomy information helps the model to avoid semantically implausible predictions.




Figure 3: Correlation analysis for location and time attributes. Best viewed in color.

5.4.3 Performance Comparison for Under-represented Species. The iWildCam2020-WILDS dataset exhibits a long-tail species distribution [30], posing challenges for accurately recognizing species that are under-represented in the training set. We compare the performance of our best-performing model (COSMO with location context) to the baseline Empirical Risk Minimization (ResNet-50) model (Table 6). We focus on examples whose labels have a maximum of 100 instances in the training set and report the overall test set accuracy for this subset of species. This selection includes species like banded palm civet, Brazilian cottontail, and leopard, all classified as vulnerable in IUCN’s list of threatened species [66]. We observe that COSMO outperforms the ERM baseline by 2.7%,




Figure 4: Plot of location GPS coordinates for training and validation splits (iWildCam2020-WILDS). The coordinates can be grouped into six clusters. Most coordinates exhibit an overlap with their respective cluster centroids at this visualization scale. Best viewed in color.

which is a 16.6% relative improvement. These findings illustrate the potential of our model to mitigate sample inefficiency in existing approaches for under-represented species by utilizing multimodal context information.

6 Discussion and Conclusion

In this work, we presented a novel framework in which the species classification task is reformulated as link prediction in a multimodal KG of species images and their diverse contextual information. This enables a unified way to leverage various forms of multimodal

context, *e.g.*, numerical, categorical, and taxonomy information associated with images for species classification in camera traps. Through our experiments, we demonstrate that our framework achieves superior out-of-distribution generalization and competitive performance with state-of-the-art for species classification on the iWildCam2020-WILDS and Snapshot Mountain Zebra datasets. Additionally, our framework exhibits improved sample efficiency in recognizing under-represented and vulnerable wildlife species.

We assume that there is a perfect linkage between these contexts and the corresponding images in the training set. However, in scenarios where such linkage is unavailable, the training procedure may introduce noise, which could lead to inferior generalization capabilities in the model. Additionally, it is important to note that the effectiveness of diverse contexts varies based on their informativeness for the given task. Interestingly, combining two or more contexts could degrade performance compared to using a single context type in some cases (Table 1). We posit that specific metadata, like location, might have a stronger regularization effect on improving generalization in species recognition tasks than other metadata. To address this, future work will involve enabling the model to assign greater importance to more informative metadata.

Furthermore, we are interested in training a foundation model for camera trap species classification across a wider spectrum of species. This model should demonstrate enhanced generalization capabilities for new camera trap setups worldwide. Additionally, we aim to integrate a broader spectrum of diverse contexts such as temperature, weather conditions, habitat, and sequence information for use with real-world camera trap deployments.

Acknowledgements

The authors would like to thank colleagues from the OSU NLP group for their constructive feedback. This research was sponsored in part by NSF OAC 2112606, NSF OAC 2118240, Cisco, and the Ohio Supercomputer Center [14].

References

- [1] Jorge A Ahumada, Eric Fegraus, Tanya Birch, Nicole Flores, Roland Kays, Timothy G O'Brien, Jonathan Palmer, Stephanie Schuttler, Jennifer Y Zhao, Walter Jetz, et al. 2020. Wildlife insights: A platform to maximize the potential of camera trap and other passive sensor wildlife data for the planet. *Environmental Conservation* 47, 1 (2020), 1–6.
- [2] Rosamund EA Almond, Monique Grooten, and T Peterson. 2020. *Living Planet Report 2020-Bending the curve of biodiversity loss*. World Wildlife Fund.
- [3] Bilal Alsallakh, Amin Jourabloo, Mao Ye, Xiaoming Liu, and Liu Ren. 2018. Do Convolutional Neural Networks Learn Class Hierarchy? *IEEE Trans. Vis. Comput. Graph.* 24, 1 (2018), 152–162. <https://doi.org/10.1109/TVCG.2017.2744683>
- [4] Elissa Aminoff, Nurit Gronau, and Moshe Bar. 2007. The parahippocampal cortex mediates spatial and nonspatial associations. *Cerebral cortex* 17, 7 (2007), 1493–1503.
- [5] Moshe Bar. 2004. Visual objects in context. *Nature Reviews Neuroscience* 5, 8 (2004), 617–629.
- [6] Suchet Bargoti and James Underwood. 2016. Image classification with orchard metadata. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 5164–5170.
- [7] Sara Beery, Elijah Cole, and Arvi Gjoka. 2020. The iWildCam 2020 Competition Dataset. arXiv:2004.10340 [cs.CV] <https://arxiv.org/abs/2004.10340>
- [8] Sara Beery, Grant van Horn, Oisin Mac Aodha, and Pietro Perona. 2019. The iWildCam 2018 Challenge Dataset. arXiv:1904.05986 [cs.CV] <https://arxiv.org/abs/1904.05986>
- [9] Sara Beery, Grant Van Horn, and Pietro Perona. 2018. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*. 456–473.
- [10] Luca Bertinetto, Romain Müller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A. Lord. 2020. Making Better Mistakes: Leveraging Class Hierarchies With Deep Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*. Computer Vision Foundation / IEEE, 12503–12512. <https://doi.org/10.1109/CVPR42600.2020.01252>
- [11] Anil Bhattacharyya. 1946. On a measure of divergence between two multinomial populations. *Sankhyā: the indian journal of statistics* (1946), 401–406.
- [12] Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5–8, 2013, Lake Tahoe, Nevada, United States*, Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger (Eds.). 2787–2795. <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>
- [13] Ludwig Bothmann, Lisa Wimmer, Omid Charrakh, Tobias Weber, Hendrik Edelhoff, Wibke Peters, Hien Nguyen, Caryl Benjamin, and Annette Menzel. 2023. Automated wildlife image classification: An active learning tool for ecological applications. *Ecol. Informatics* 77 (2023), 102231. <https://doi.org/10.1016/J.ECOINF.2023.102231>
- [14] Ohio Supercomputer Center. 1987. Ohio Supercomputer Center. <http://osc.edu/ark:/19495/f5s1ph73>
- [15] Sanxing Chen, Xiaodong Liu, Jianfeng Gao, Jian Jiao, Ruofei Zhang, and Yangfeng Ji. 2021. HittER: Hierarchical Transformers for Knowledge Graph Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 10395–10407. <https://doi.org/10.18653/v1/2021.emnlp-main.812>
- [16] Grace Chu, Brian Potetz, Weijun Wang, Andrew Howard, Yang Song, Fernando Brucher, Thomas Leung, and Hartwig Adam. 2019. Geo-Aware Networks for Fine-Grained Recognition. In *2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27–28, 2019*. IEEE, 247–254. <https://doi.org/10.1109/ICCVW.2019.00033>
- [17] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Thirty-second AAAI conference on artificial intelligence*.
- [18] Qishuai Diao, Yi Jiang, Bin Wen, Jia Sun, and Zehuan Yuan. 2022. MetaFormer: A Unified Meta Framework for Fine-Grained Recognition. arXiv:2203.02751 [cs.CV] <https://arxiv.org/abs/2203.02751>
- [19] Sandra Diaz, Josef Settele, Eduardo S Brondizio, Hien T Ngo, John Agard, Almut Arneth, Patricia Balvanera, Kate A Brauman, Stuart HM Butchart, Kai MA Chan, et al. 2019. Pervasive human-driven decline of life on Earth points to the need for transformative change. *Science* 366, 6471 (2019), eaax3100.
- [20] Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 601–610.
- [21] Jeffrey S Ellen, Casey A Graff, and Mark D Ohman. 2019. Improving plankton image classification using context metadata. *Limnology and Oceanography*:

- Methods* 17, 8 (2019), 439–461.
- [22] Alberto García-Durán and Mathias Niepert. 2018. KBLrn: End-to-End Learning of Knowledge Base Representations with Latent, Relational, and Numerical Features. In *Conference on Uncertainty in Artificial Intelligence*. <http://auai.org/uai2018/proceedings/papers/149.pdf>
 - [23] Paul Glover-Kapfer, Carolina A Soto-Navarro, and Oliver R Wearn. 2019. Camera-trapping version 3.0: current constraints and future priorities for development. *Remote Sensing in Ecology and Conservation* 5, 3 (2019), 209–223.
 - [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016*. IEEE Computer Society, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
 - [25] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. 2018. The INaturalist Species Classification and Detection Dataset. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. Computer Vision Foundation / IEEE Computer Society, 8769–8778. <https://doi.org/10.1109/CVPR.2018.00914>
 - [26] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. 2018. Does Distributionally Robust Supervised Learning Give Robust Classifiers?. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10–15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 2034–2042. <http://proceedings.mlr.press/v80/hu18a.html>
 - [27] Mirantha Jayatilaka, Tingting Mu, and Uli Sattler. 2021. Ontology-based n-ball Concept Embeddings Informing Few-shot Image Classification. In *Machine Learning with Symbolic Methods and Knowledge Graphs co-located with European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD 2021), Virtual, September 17, 2021 (CEUR Workshop Proceedings, Vol. 2997)*, Mehdi Ali, Paul Groth, Pascal Hitzler, Jens Lehmann, Heiko Paulheim, Achim Rettinger, Harald Sack, Afshin Sadeghi, and Volker Tresp (Eds.). CEUR-WS.org. <https://ceur-ws.org/Vol-2997/paper1.pdf>
 - [28] Justin Johnson, Lamberto Ballan, and Li Fei-Fei. 2015. Love Thy Neighbors: Image Annotation by Exploiting Image Metadata. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7–13, 2015*. IEEE Computer Society, 4624–4632. <https://doi.org/10.1109/ICCV.2015.525>
 - [29] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
 - [30] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran S. Haque, Sara M. Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 5637–5664. <http://proceedings.mlr.press/v139/koh21a.html>
 - [31] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *Int. J. Comput. Vis.* 123, 1 (2017), 32–73. <https://doi.org/10.1007/S11263-016-0981-7>
 - [32] Wen Li, Li Niu, and Dong Xu. 2014. Exploiting Privileged Information from Web Data for Image Categorization. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V (Lecture Notes in Computer Science, Vol. 8693)*, David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer, 437–452. https://doi.org/10.1007/978-3-319-10602-1_29
 - [33] Xinhang Li, Xiangyu Zhao, Jiaxing Xu, Yong Zhang, and Chunxiao Xing. 2023. IMF: Interactive Multimodal Fusion Model for Link Prediction. In *Proceedings of the ACM Web Conference 2023*. 2572–2580.
 - [34] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 2181–2187. <http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9571>
 - [35] Chengjiang Long, Roddy Collins, Eran Swears, and Anthony Hoogs. 2019. Deep Neural Networks in Fully Connected CRF for Image Labeling with Social Network Metadata. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7–11, 2019*. IEEE, 1607–1615. <https://doi.org/10.1109/WACV.2019.00176>
 - [36] Kenneth Marino, Ruslan Salakhutdinov, and Abhinav Gupta. 2017. The More You Know: Using Knowledge Graphs for Image Classification. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*. IEEE Computer Society, 20–28. <https://doi.org/10.1109/CVPR.2017.10>
 - [37] Sean L Maxwell, Richard A Fuller, Thomas M Brooks, and James EM Watson. 2016. Biodiversity: The ravages of guns, nets and bulldozers. *Nature* 536, 7615 (2016), 143–145.
 - [38] Julian J. McAuley and Jure Leskovec. 2012. Image Labeling on a Network: Using Social-Network Metadata for Image Classification. In *Computer Vision - ECCV 2012 - 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part IV (Lecture Notes in Computer Science, Vol. 7575)*, Andrew W. Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid (Eds.). Springer, 828–841. https://doi.org/10.1007/978-3-642-33765-9_59
 - [39] Zhongqi Miao, Kaitlyn M Gaynor, Jiayun Wang, Ziwei Liu, Oliver Muellerklein, Mohammad Sadegh Norouzzadeh, Alex McInturff, Rauri CK Bowie, Ran Nathan, Stella X Yu, et al. 2019. Insights and approaches using deep learning to classify wildlife. *Scientific reports* 9, 1 (2019), 8137.
 - [40] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.
 - [41] Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2018. A Novel Embedding Model for Knowledge Base Completion Based on Convolutional Neural Network. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 327–333. <https://doi.org/10.18653/v1/N18-2053>
 - [42] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, Lise Getoor and Tobias Scheffer (Eds.). Omnipress, 809–816. https://icml.cc/2011/papers/438_icmlpaper.pdf
 - [43] Mohammad Sadegh Norouzzadeh, Anh Nguyen, Margaret Kosmala, Alexandra Swanson, Meredith S Palmer, Craig Packer, and Jeff Clune. 2018. Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences* 115, 25 (2018), E5716–E5725.
 - [44] Allan F O’Connell, James D Nichols, and K Ullas Karanth. 2011. *Camera traps in animal ecology: methods and analyses*. Vol. 271. Springer.
 - [45] Aude Oliva and Antonio Torralba. 2007. The role of context in object recognition. *Trends in cognitive sciences* 11, 12 (2007), 520–527.
 - [46] OpenTreeofLife, Karen A. Cranston, Benjamin Redelings, Luna Luisa Sanchez Reyes, Jim Allman, Emily Jane McTavish, and Mark T. Holder. 2019. Open Tree of Life Taxonomy. <https://doi.org/10.5281/zenodo.3937751>
 - [47] Vardaan Pahuja, Boshi Wang, Hugo Latapie, Jayanth Srinivasa, and Yu Su. 2023. A Retrieve-and-Read Framework for Knowledge Graph Link Prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (Birmingham, United Kingdom) (CIKM ’23)*. Association for Computing Machinery, New York, NY, USA, 1992–2002. <https://doi.org/10.1145/3583780.3614769>
 - [48] Lain E Pardo, Sara Bombaci, Sarah E Huebner, Michael J Somers, Herve Fritz, Colleen Downs, Abby Guthmann, Robyn S Hetem, Mark Keith, Aliza le Roux, et al. 2021. Snapshot Safari: A large-scale collaborative to monitor Africa’s remarkable biodiversity. *South African Journal of Science* 117, 1–2 (2021), 1–4.
 - [49] Pouya Pezeshkpour, Liyan Chen, and Sameer Singh. 2018. Embedding Multimodal Relational Data for Knowledge Base Completion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 3208–3218. <https://doi.org/10.18653/v1/D18-1359>
 - [50] Qi Qi, Yi Xu, Wotao Yin, Rong Jin, and Tianbao Yang. 2023. Attentional-Biased Stochastic Gradient Descent. *Transactions on Machine Learning Research* (2023). <https://openreview.net/forum?id=B0WYwVVA2r>
 - [51] Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. 2020. You CAN Teach an Old Dog New Tricks! On Training Knowledge Graph Embeddings. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BkxSmlBFvr>
 - [52] Mohammad Sadegh Norouzzadeh, Dan Morris, Sara Beery, Neel Joshi, Nebojsa Jojic, and Jeff Clune. 2020. A Deep Active Learning System for Species Identification and Counting in Camera Trap Images. *Methods in Ecology and Evolution* 12, 1 (December 2020), 150–161. <https://www.microsoft.com/en-us/research/publication/a-deep-active-learning-system-for-species-identification-and-counting-in-camera-trap-images/>
 - [53] Apoorv Saxena, Adrian Kochsiek, and Rainer Gemulla. 2022. Sequence-to-Sequence Knowledge Graph Completion and Question Answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 2814–2828. <https://doi.org/10.18653/v1/2022.acl-long.201>
 - [54] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*. Springer, 593–607.
 - [55] Stefan Schneider, Saul Greenberg, Graham W Taylor, and Stefan C Kremer. 2020. Three critical factors affecting automated image species recognition performance for camera traps. *Ecology and evolution* 10, 7 (2020), 3503–3517.
 - [56] Stefan Schneider, Graham W. Taylor, and Stefan C. Kremer. 2018. Deep Learning Object Detection Methods for Ecological Camera Trap Data. In *15th Conference*

- on *Computer and Robot Vision, CRV 2018, Toronto, ON, Canada, May 8-10, 2018*. IEEE Computer Society, 321–328. <https://doi.org/10.1109/CRV.2018.00052>
- [57] Hatem Mousselly Sergieh, Teresa Botschen, Iryna Gurevych, and Stefan Roth. 2018. A Multimodal Translation-Based Approach for Knowledge Graph Representation Learning. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics, *SEM@NAACL-HLT 2018, New Orleans, Louisiana, USA, June 5-6, 2018*, Malvina Nissim, Jonathan Berant, and Alessandro Lenci (Eds.). Association for Computational Linguistics, 225–234. <https://doi.org/10.18653/v1/s18-2027>
- [58] Yuge Shi, Jeffrey Seely, Philip Torr, Siddharth N, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. 2022. Gradient Matching for Domain Generalization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=vDwBW49HmO>
- [59] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1409.1556>
- [60] Samuel Stevens, Jiaman Wu, Matthew J Thompson, Elizabeth G Campolongo, Chan Hee Song, David Edward Carlyn, Li Dong, Wasila M Dahdul, Charles Stewart, Tanya Berger-Wolf, Wei-Lun Chao, and Yu Su. 2024. BioCLIP: A Vision Foundation Model for the Tree of Life. In *2024 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 17-21, 2024*. Computer Vision Foundation / IEEE Computer Society, 19412–19424. https://openaccess.thecvf.com/content/CVPR2024/html/Stevens_BioCLIP_A_Vision_Foundation_Model_for_the_Tree_of_Life_CVPR_2024_paper.html
- [61] Baochen Sun and Kate Saenko. 2016. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In *Computer Vision - ECCV 2016 Workshops - Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 9915)*, Gang Hua and Hervé Jégou (Eds.). 443–450. https://doi.org/10.1007/978-3-319-49409-8_35
- [62] Zhiqing Sun, Shikhar Vashishth, Soumya Sanyal, Partha Talukdar, and Yiming Yang. 2020. A Re-evaluation of Knowledge Graph Completion Methods. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 5516–5522. <https://doi.org/10.18653/v1/2020.acl-main.489>
- [63] Michael A Tabak, Mohammad S Norouzzadeh, David W Wolfson, Erica J Newton, Raoul K Boughton, Jacob S Ivan, Eric A Odell, Eric S Newkirk, Reesa Y Conrey, Jennifer Stenglein, et al. 2020. Improving the accessibility and transferability of machine learning algorithms for identification of animals in camera trap images: MLWIC2. *Ecology and evolution* 10, 19 (2020), 10374–10383.
- [64] Michael A Tabak, Mohammad S Norouzzadeh, David W Wolfson, Steven J Sweeney, Kurt C VerCauteren, Nathan P Snow, Joseph M Halseth, Paul A Di Salvo, Jesse S Lewis, Michael D White, et al. 2019. Machine learning to classify animal species in camera trap images: Applications in ecology. *Methods in Ecology and Evolution* 10, 4 (2019), 585–590.
- [65] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. 2020. Composition-based Multi-Relational Graph Convolutional Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. https://openreview.net/forum?id=BylA_C4tPr
- [66] Jean-Christophe Vié, Craig Hilton-Taylor, Caroline Pollock, James Ragle, Jane Smart, Simon N Stuart, and Rashila Tong. 2009. The IUCN Red List: a key conservation tool. *Wildlife in a changing world—An analysis of the 2008 IUCN Red List of Threatened Species* (2009), 1.
- [67] OR Wearn and P Glover-Kapfer. 2017. Camera-trapping for conservation: a guide to best-practices. *WWF conservation technology series* 1, 1 (2017), 181.
- [68] Ben G Weinstein. 2018. A computer vision for animal ecology. *Journal of Animal Ecology* 87, 3 (2018), 533–545.
- [69] Xander Wilcke, Peter Bloem, Victor de Boer, and Rein van 't Veer. 2021. End-to-End Learning on Multimodal Knowledge Graphs. *Semantic Web - Interoperability, Usability, Applicability an IOS Press Journal* (2021). <https://www.semantic-web-journal.net/content/end-end-learning-multimodal-knowledge-graphs>
- [70] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6575>
- [71] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. KG-BERT: BERT for Knowledge Graph Completion. arXiv:1909.03193 [cs.CL] <https://arxiv.org/abs/1909.03193>
- [72] Donghan Yu, Yiming Yang, Ruohong Zhang, and Yuexin Wu. 2021. Knowledge Embedding Based Graph Convolutional Network. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila Zia (Eds.). ACM / IW3C2, 1619–1628. <https://doi.org/10.1145/3442381.3449925>
- [73] Shu Zhang, Ran Xu, Caiming Xiong, and Chetan Ramaiah. 2022. Use All The Labels: A Hierarchical Multi-Label Contrastive Learning Framework. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. IEEE, 16639–16648. <https://doi.org/10.1109/CVPR52688.2022.01616>




Figure C.1: Species probabilities conditioned on day/night for the 10 most frequent species in the training set (iWildCam2020-WILDS). Animal species demonstrate distinct temporal preferences for their daily activities, as evidenced by the contrasting probabilities observed during day and night.

Supplementary Material

A Datasets

The iWildCam2020-WILDS dataset [30] and Snapshot Mountain Zebra dataset [48] are released under the Community Data License Agreement (CDLA). The Open Tree Taxonomy [46] is licensed under Creative Commons Attribution 1.0 Generic license. Both licenses permit their use for academic research in their original form. The iWildCam2020-WILDS label space consists of 182 species. Out of these 182 species, 155 species have support in the OTT taxonomy.

B Experiment Details

Table B.1 and Table B.2 show the training time of COSMO under different settings for iWildCam2020-WILDS and Snapshot Mountain

Zebra, respectively. We performed all experiments using a single Nvidia RTX A6000 GPU.

Table B.1: Training time for different ablations of COSMO (iWildCam2020-WILDS). All experiments use a single Nvidia RTX A6000 GPU.

Multi-modality			Running time (hrs.)
Taxonomy	Location	Time	
	–		3.9
✓			3.9
	✓		6.3
		✓	7.5
✓	✓		6.3
✓		✓	7.5
	✓	✓	9.9
✓	✓	✓	9.9

Table B.2: Training time for different ablations of COSMO (Snapshot Mountain Zebra). All experiments use a single Nvidia RTX A6000 GPU.

Multi-modality		Running time (hrs.)
Taxonomy	Time	
	–	4.2
✓		4.2
	✓	7.0
✓	✓	7.0

C Time Attribute Analysis

We analyzed the time attributes associated with the images of the 10 most frequent species in the training set of iWildCam2020-WILDS dataset (Figure C.1). We observed a significant difference in the species distributions between day and night⁴. This indicates that temporal information plays a crucial role in recognizing these species.

⁴In this analysis, we define the time duration between 5 A.M. and 7 P.M. local time as daytime.