Multi-modal Conditional Bounding Box Regression for Music Score Following

Florian Henkel¹, Gerhard Widmer^{1,2} ¹Institute of Computational Perception & ²LIT Artificial Intelligence Lab Johannes Kepler University Linz, Austria florian.henkel@jku.at

Abstract—This paper addresses the problem of sheet-imagebased on-line audio-to-score alignment also known as score following. Drawing inspiration from object detection, a conditional neural network architecture is proposed that directly predicts x,y coordinates of the matching positions in a complete score sheet image at each point in time for a given musical performance.

Experiments are conducted on a synthetic polyphonic piano benchmark dataset and the new method is compared to several existing approaches from the literature for sheet-image-based score following as well as an Optical Music Recognition baseline.

The proposed approach achieves new state-of-the-art results and furthermore significantly improves the alignment performance on a set of real-world piano recordings by applying Impulse Responses as a data augmentation technique.

Index Terms—audio-to-score alignment, score following, conditional object detection, multi-modal deep learning

I. INTRODUCTION

Audio-to-score alignment is a well known task in music signal processing [1]–[5], where a musical performance and a corresponding score have to be synchronized along the time axis. This can be approached as an offline task when audio and score are available upfront, or in an on-line fashion if the complete audio is not immediately accessible but incrementally processed. This is also referred to as real-time alignment or score following. Score following has many applications in music, such as automatic page turning for musicians [6] or synchronised information display in the concert hall [7].

While traditional score following methods usually rely on computer-readable scores [8]–[11], e. g. MIDI or MusicXML, recent deep learning based approaches try to avoid the need of such score representations by directly following along in a sheet image [12]–[15]. Initially, methods falling in the latter category required the score image to be available in an unrolled form: the staves have to be detected on the page, cut out and stitched together to a long sequence. To avoid the need for an external system for this sub-problem, [15] proposed to treat score following as a *referring image segmentation* task, which for the first time allowed to follow musical performances in full-page sheet images without any pre-processing.

One drawback of this approach is the predicted segmentation mask, which assigns a probability to each pixel in the sheet image indicating whether it corresponds to the given audio. Apart from being counter-intuitive, since there is no clear musical correspondence between a pixel by itself and the audio, it also allows the system to predict multiple regions on the sheet image (see Figure 1 in [15]).

To overcome this, we propose a conditional object detection architecture based on the family of YOLO object detectors [16], and subsequently treat score following as a bounding box regression instead of a segmentation task. By incorporating a confidence score (to be described in Section III-A) our system offers a clear distinction between different predicted positions on the sheet image. We empirically show that our proposed method outperforms existing approaches on MSMD, a synthetic polyphonic piano benchmark dataset [17].

A further problem in previous work is the performance gap between the synthetic audio the models are trained on compared to real-world piano recordings. To alleviate this, we show that *Impulse Responses* (IRs) used for data-augmentation are a simple yet effective way to allow the network to generalize better across various recording scenarios.¹

II. RELATED PREVIOUS WORK

Dorfer et al. [12] propose to treat score following as a *one-dimensional localization task*, where a small sheet image excerpt is discretized into k bins and the most likely bin matching a given short audio excerpt has to be predicted. While our proposed approach can also be seen as a localization task, we do not rely on sheet image excerpts and discretization, but directly predict the matching position for a musical performance in the complete sheet image (see Figure 1).

In [13] and [14], score following is formulated as a *rein-forcement learning problem*. Agents are trained to follow along a musical performance in an unrolled score image by adapting their reading speed, i. e., how fast they want to move forward or backward in the score.

In [18], a conditional U-Net architecture is proposed that predicts a *segmentation mask* for all matching positions in a complete score sheet image for monophonic piano music. This was later extended in [15] to a full score-following setup incorporating long temporal context on polyphonic piano music. As a conditioning mechanism the *Feature-wise linear modulation* (FiLM) [19] layer is used, which we also apply in our proposed conditional object detection architecture.

In terms of bounding box prediction in images, current object detection systems heavily rely on deep learning and can

¹Code and data will be made available on-line: https://github.com/CPJKU/ cyolo_score_following



Fig. 1. Score following as a bounding box regression task. For each audio frame, the model is asked to output the most likely position in the sheet image that matches the audio up until the current point in time. The colored bars in the spectrogram correspond to the bounding boxes in the sheet image.

be broadly categorized into one-stage detectors, e. g. YOLO [16], and two-stage detectors, e. g. Faster R-CNN [20]. YOLObased object detectors are known for their fast inference and real-time capabilities, which is also a requirement for on-line score following and the reason we choose it for our work.

III. CONDITIONAL OBJECT DETECTION FOR MUSIC SCORE FOLLOWING

The goal in object detection is to predict bounding boxes for all objects in an image that belong to classes the model has seen during training. In contrast, we have a *conditional* setup: we want to query the model for a particular object instance, i.e., instead of predicting bounding boxes for all objects in an image, the model should only predict the bounding box for a specific object based on a query.² Similar to referring image segmentation [21], this query could be some language expression referring to the desired object. In our work, instead of language expressions, we use an audio signal as the conditioning information, similar to [15]. More specifically, we will query the object detection model repeatedly for every incoming audio frame (roughly every 50ms) asking for the position in the sheet image that matches the incoming music up to the current time step (cf. Figure 1). This is implemented as a bounding box regression task, i. e., the task of the network will be to predict the x, y center-coordinates as well as the width and height of the bounding box that matches the query. In this particular work, the height depends on the height of the corresponding staff and the width is arbitrarily chosen to be 30 pixels wide in the training data annotations.

²To clarify: what is to be enclosed in a bounding box, in our case, is not an 'object' proper but a position in the score sheet, with the bounding box indicating the notes or chord corresponding to the notes that have just started in the audio (see Figure 1). Desired width and height of such boxes are implicitly defined by our training data. For the same reason, we will evaluate our score follower (in Section V) by measuring the precision with which it tracks note onsets, rather than by bounding-box overlap between prediction and ground truth; after all, this is what will matter in score following.

TABLE I

The conditional YOLO architecture consisting of several Downscale and Upscale blocks as depicted in Figure 2. Upscale blocks concatenate the input from a previous layer given in parentheses, e.g., layer 6 takes the output of layer 4 as additional input. FiLM indicates that the conditional layer is active in these blocks whereas it is bypassed in the others.

Conditional YOLO						
Layer	Module	Channels	Output Size			
1	Focus	16	208×208			
2	Downscale	32	104×104			
3	Downscale	64	52×52			
4	Downscale - FiLM	128	26×26			
5	Downscale - FiLM	128	13×13			
6	Upscale(4) - FiLM	128	26×26			
7	Upscale(3) - FiLM	128	52×52			
8	Detection	15	52×52			



Fig. 2. The main building blocks of the architecture shown in Table I. The *Downscale* block applies a 3×3 convolution with a stride of 2, thus halving the spatial dimension of the input feature maps. The FiLM layer in both the *Downscale* and *Upscale* block is optional and, if not indicated otherwise, simply bypassed, i. e., the output of the convolutional layer is directly fed to the activation function. The *Upscale* block applies nearest neighbour up-sampling by a factor of 2 and performs a channel-wise concatenation with the feature maps from a given previous layer. The *Bottleneck* block has a residual element-wise addition and reduces the number of parameters by performing the 3×3 convolution with only half the number of output channels, e. g., $64 \times 1 \times 1$, $64 \times 3 \times 3$ and $128 \times 1 \times 1$. If input and output do not share the same number of channels an additional 1×1 convolution followed by layer normalization is applied in the residual branch before the element-wise addition.

A. Conditional YOLO

Our proposed model belongs to the family of YOLO object detectors [16] and consists of residual *Downscale* and *Upscale* blocks incorporating *Feature-wise linear modulation* (FiLM) [19] layers at several positions (cf. Table I and Figure 2).

The FiLM layers are used to inject external information (given as a conditioning vector z) into the network, by performing the following computation on feature maps x

$$f_{\rm FiLM}(\mathbf{x}) = \mathbf{s}(\mathbf{z}) \cdot \mathbf{x} + \mathbf{t}(\mathbf{z}), \tag{1}$$

with $s(\cdot)$ and $t(\cdot)$ being learned linear functions that scale and translate the input x based on z. The computation of z is described in Section III-B. For more information we refer the reader to [15].

As shown in Table I, the first layer in the network is a so-called *Focus* or *SpaceToDepth* [22] layer which splits a $c \times w \times h$ image (in our case of shape $1 \times 416 \times 416$) into

a $4 * c \times \frac{w}{2} \times \frac{h}{2}$ (4 × 208 × 208) dimensional image with the purpose of reducing the image size and improving the overall computation speed. This is followed directly by a 16 × 3 × 3 convolution, layer normalization [23] and ELU activation [24]. Each *Downscale* block reduces the input size by half, thus creating feature maps of different resolutions (104 × 104, 52 × 52, 26 × 26, 13 × 13). In the *Upscale* blocks, these feature maps are up-sampled, concatenated, and subsequently used in the *Detection* layer to create the final predictions.

The predictions are partly based on anchor points for which the *Detection* layer predicts offsets to produce the final width and height of the bounding box. Given that we consider uniform objects of a fixed width and similar height that depend on the staves, we only use three anchor points (11, 26), (11, 34), (11, 45). 11 indicates the anchor-width and 26, 34, 45 the anchor-heights for a down-scaled sheet image (cf. Section IV). These are determined using k-means clustering of the training set bounding box [25].

For each spatial position of the 52×52 output grid and for each anchor point, the *Detection* layer produces 5 outputs. The first two are offsets for the x, y center coordinates of the bounding boxes. Based on these and the spatial position in the grid, the coordinates in the original input image will be computed. The next two outputs correspond to the aforementioned width and height offsets and the last output is an objectness or confidence score predicting the *Intersection over Union* (IoU) for a bounding box. Colloquially, this tells us how well the predicted box fits the ground truth object. During inference, this score is used to filter the most likely position in the sheet image that matches the given audio.

B. Audio Encoding

The audio waveform is sampled with 22.05 kHz and subjected to an STFT with a Hann window of size 2048 and a hop size of 1102 (approx. 20 frames per second). Each frame is further transformed using a logarithmic filterbank processing frequencies between 60 Hz and 6 kHz. The resulting output has 78 log-frequency bins. Given the on-the-fly data augmentation during training to be described in Section IV-A, we cannot normalize the input spectrograms to zero mean and unit variance beforehand as in [15]. To still allow for input normalization, we add a batch normalization layer [26] in front of our encoding network that computes mean and standard deviation parameters for each frequency band similar to [27]. To encode the spectrogram, we apply the same CNN encoder as in [15], projecting 40 audio frames to a 32 dimensional vector x. To encode the larger temporal audio context, we apply an LSTM layer with 64 hidden units [28]. Its hidden state is updated every 40 audio frames using x. The final conditioning vector z that will be used in the FiLM layers is computed as

$$\mathbf{z} = f([\mathbf{h}; \mathbf{x}]), \tag{2}$$

where **h** is the 64 dimensional hidden state vector of the LSTM, $[\mathbf{h}; \mathbf{x}]$ indicates concatenation of **h** and **x**, and *f* is a fully connected layer of size 128 with layer normalization and ELU activation. Thus, the final dimensionality of **z** is 128.

IV. EXPERIMENTAL SETUP

We conduct our experiments using the Multi-modal Sheet Music Dataset (MSMD) [17], a piano music dataset which offers alignments between notehead positions in the sheet image and a synthetic MIDI score. For training, this MIDI is rendered to an audio performance using a publicly available piano sound-font with Fluidsynth.³ As in [15], the dataset is split into 945 train, 28 validation and 125 test pages. The sheet image pages with an initial shape of 1181×835 are zero-padded to a squared size and down-scaled to 416×416 . This still offers a high enough resolution such that humans are able to detect details in the sheet image.

The *optimization objective* consists of two parts, a meansquared-error loss for the bounding boxes as well as a logistic regression loss for predicting the aforementioned objectness score, which should reflect the IoU score of a predicted bounding box, i. e., a predicted box with a perfect overlap to the ground truth should have a score of 1 [25].

For training, we use the Adam optimizer with corrected weight decay regularization [29], and a weight decay coefficient of $5e^{-4}$. Weight decay is only applied to the weight parameters of convolutional, recurrent and linear layers, but not to normalization layers and bias parameters [30]. The weights of the network are initialized orthogonally [31] and the biases are set to zero except for the forget gate bias in the LSTM layer which is set to 1 [32]. We anneal the learning rate over the course of 50 epochs starting from $5e^{-4}$ to $5e^{-6}$ using a cosine annealing scheme [33]. The best models in terms of the validation loss are selected for evaluation on the test set. To avoid exploding gradients, we apply gradient clipping with a maximum norm of 0.1 to the parameters of the recurrent layer in the audio encoder.

We perform data augmentation in the sheet image domain by randomly shifting the score along the x, y axis and additionally augment the audio by changing the tempo of the performance with a random factor between 0.5 and 2.

A. Impulse Response Data Augmentation

One issue with existing approaches to sheet-image-based score following is their lack of generalization to real-world audio [14], [15]. This is attributed to the synthetic data these models are trained on, due to the shortage of alignments between real-world audio recordings and notehead positions in sheet music. To alleviate this, we apply *Impulse Responses* (IRs) as an audio data augmentation technique. IRs allow us to model different recording conditions in the form of microphones and room characteristics. We use more than 500 freely available IRs collected from OpenAIRLib⁴ and MicIRP.⁵ The IR signal is convolved with the audio on-the-fly during training, such that in each epoch the model encounters different audio scenarios, which should allow for a more robust audio encoding model. To measure the effect of this technique, we train our model with and without IR augmentation.

³https://www.fluidsynth.org

⁴https://openairlib.net

⁵https://micirp.blogspot.com

B. Baselines

As baselines we choose three methods for sheet-imagebased score following: the supervised localization network referred to as MM-Loc [12], the best performing reinforcement learning (RL) agent from [14], and the conditional segmentation network [15], which we refer to as CUNet. Additionally, a more traditional method based on optical music recognition (OMR) is considered that first performs OMR on a given sheet image to create a MIDI score representation. The MIDI is then rendered to audio and on-line dynamic time warping (ODTW) [34] is applied to perform audio-to-audio alignment.

V. RESULTS

Table II summarizes the results across different settings. As an evaluation metric, we choose the ratio of tracked onsets below certain error thresholds ranging from 0.05 to 5 seconds similar to [15].⁶ To arrive at this metric, we transform the predicted bounding boxes in the sheet image back to the corresponding timestep position in the audio using the ground truth annotations. The absolute difference between the predicted timestep (transformed from the bounding box) and the ground truth note onset then yields the tracking error. In the first section of the table, we evaluate all approaches on the full MSMD test split consisting of 125 sheet-image pages. This is a completely synthetic setup, meaning that the audio is rendered from MIDI with the same piano sound-font seen during training. We observe that our non IR-augmented model (CYOLO) already outperforms all previous methods across all error thresholds. Additionally incorporating IR-augmentation (CYOLO-IR) slightly improves the tracking performance, potentially by making it harder to overfit to the training data. As a notable difference to the referring image segmentation method (CUNet), our model not only tracks notes more accurately but is overall able to track more notes below the 5 second error threshold. We attribute this to our task formulation that allows us to directly predict positions in the sheet image instead of inferring them from segmented regions.

Starting with section II of the table, we now limit the evaluation to a subset of 25 pages from the MSMD test set for which we also have real piano recordings that were aligned to the sheet images. This will allow us to measure how much expressive performance deviations, played ornaments, errors, and different audio conditions affect the score followers. As a baseline, section II shows the results on the synthetic audios for just this subset of pieces.

In the remaining three sections (III a-c), we have the same 25 pages from the test set, with (a) the recorded *performance* MIDI synthesized by the sound-font used during training, (b) the audio recorded from the direct output of a Yamaha AvantGrand N2 hybrid piano, and (c) the audio captured with a microphone in an office room.

(a) allows us to observe the effect of expressive playing in contrast to a deadpan MIDI performance, while still having a

TABLE II

Comparison of our proposed methods to several approaches as described in Section IV. We report the ratio of tracked onsets below certain error thresholds from 0.05 to 5 seconds. The best result for each threshold is marked bold.

Err. [sec]	≤ 0.05	≤ 0.10	≤ 0.50	≤ 1.00	≤ 5.00			
I MSMD Full Test Set (Synthetic)								
OMR [14]	0.447	0.519	0.760	0.850	0.974			
MM-Loc [12]	0.446	0.492	0.822	0.860	0.920			
RL [14]	0.409	0.433	0.797	0.878	0.972			
CUNet [15]	0.733	0.747	0.852	0.885	0.937			
CYOLO	0.815	0.826	0.865	0.890	0.979			
CYOLO-IR	0.830	0.842	0.885	0.909	0.984			
II MSMD Test Subset (Synthetic)								
OMR [14]	0.371	0.461	0.749	0.868	0.996			
MM-Loc [12]	0.416	0.442	0.776	0.799	0.903			
RL [14]	0.365	0.382	0.729	0.798	0.965			
CUNet [15]	0.698	0.706	0.806	0.824	0.891			
CYOLO	0.768	0.777	0.829	0.846	0.979			
CYOLO-IR	0.795	0.803	0.863	0.880	0.990			
III (a) Performance MIDI Synthesized								
OMR [14]	0.289	0.398	0.717	0.834	0.988			
MM-Loc [12]	0.472	0.490	0.832	0.861	0.960			
RL [14]	0.234	0.248	0.545	0.640	0.812			
CUNet [15]	0.565	0.581	0.809	0.844	0.901			
CYOLO	0.661	0.676	0.804	0.840	0.962			
CYOLO-IR	0.759	0.774	0.856	0.882	0.987			
III (b) Direct Out								
OMR [14]	0.226	0.330	0.703	0.839	0.993			
MM-Loc [12]	0.338	0.354	0.597	0.634	0.753			
RL [14]	0.277	0.291	0.607	0.733	0.955			
CUNet [15]	0.400	0.416	0.642	0.693	0.811			
CYOLO	0.584	0.599	0.769	0.817	0.973			
CYOLO-IR	0.642	0.656	0.772	0.806	0.965			
III (c) Room Recording								
OMR [14]	0.226	0.322	0.702	0.827	0.974			
MM-Loc [12]	0.207	0.243	0.541	0.573	0.702			
RL [14]	0.192	0.206	0.466	0.587	0.891			
CUNet [15]	0.094	0.105	0.215	0.262	0.443			
CYOLO	0.306	0.323	0.460	0.510	0.763			
CYOLO-IR	0.563	0.581	0.712	0.749	0.919			

similar audio condition as during training. Indeed, we see some slight deterioration relative to II, overall. Interestingly, even though we are still in the synthesized piano sound domain, the IR audio augmentation again seems to help – this is likely due to the loudness variations that come with the real performances but did not appear in the original synthesized audios.

In (b) we see the first results for different playing as well as audio conditions. Again, our method performs best across the three lowest error thresholds and using IR-augmentation further improves the results. Apart from the OMR approach, we outperform all other sheet-image-based methods.

(c) offers the most difficult challenge as the audio conditions are particularly different to what was seen during training,

⁶Note that eventually, realistic evaluation criteria should reflect the targeted application. While accompaniment systems require very precise results (i. e., low error thresholds), a page turner can get away with lower accuracy.

due to having been recorded in an office environment, including background noise as well as reverberation effects from the room. In this setting we clearly observe how the IRaugmentation considerably improves the results. Our method outperforms the sheet-image-based approaches across all, and the OMR-based approach for the three lowest error thresholds. There is still a performance gap to the strictly synthetic setting, however our method offers a big step towards closing it. We assume that without the use of additional (real-world) data we will probably not be able to further reduce this gap.

VI. CONCLUSION

We proposed a new architecture for sheet-image-based score following, inspired by an object detection setup. The method compares favorably to existing methods from the literature and significantly improves the performance for real-world piano recordings. Besides, our new approach also promises to be easily extendable. In future work, we plan to not only predict the most likely note-level position, but also the corresponding bar and system. We argue, that this could lead to additional stability for our network and furthermore allow us to leverage more data, where no fine-granular note-level alignment between audio and sheet music is available. This would make it easier to use scanned or photographed scores, and thus make it possible to improve and evaluate the generalization capabilities not only in the audio but also in the sheet-image domain.

ACKNOWLEDGMENT

This work was supported by the European Research Council (ERC) under the EU's Horizon 2020 research and innovation program (grant number 670035 "Con Espressione").

REFERENCES

- N. Montecchio and A. Cont, "A unified approach to real time audioto-score and audio-to-audio alignment using sequential Montecarlo inference techniques," in *Proc. of the IEEE Int. Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.
- [2] A. Cont, "Realtime Audio to Score Alignment for Polyphonic Music Instruments using Sparse Non-Negative Constraints and Hierarchical HMMS," in Proc. of the IEEE Int. Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2006.
- [3] T. Otsuka, K. Nakadai, T. Takahashi, T. Ogata, and H. Okuno, "Realtime audio-to-score alignment using particle filter for coplayer music robots," *EURASIP Journal on Advances in Signal Processing*, vol. 2011, 2011.
- [4] R. Agrawal and S. Dixon, "Learning Frame Similarity using Siamese networks for Audio-to-Score Alignment," in *Proc. of the 28th Eur. Signal Processing Conference (EUSIPCO)*, 2020.
- [5] C. Joder, S. Essid, and G. Richard, "A Conditional Random Field Framework for Robust and Scalable Audio-to-Score Matching," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 19, no. 8, 2011.
- [6] A. Arzt, G. Widmer, and S. Dixon, "Automatic Page Turning for Musicians via Real-Time Machine Listening," in Proc. of the 18th Eur. Conference on Artificial Intelligence (ECAI), 2008.
- [7] A. Arzt, H. Frostel, T. Gadermaier, M. Gasser, M. Grachten, and G. Widmer, "Artificial Intelligence in the Concertgebouw," in *Proc. of* the 24th Int. Joint Conference on Artificial Intelligence (IJCAI), 2015.
- [8] A. Arzt, "Flexible and Robust Music Tracking," Ph.D. dissertation, Johannes Kepler University Linz, 2016.
- [9] E. Nakamura, P. Cuvillier, A. Cont, N. Ono, and S. Sagayama, "Autoregressive Hidden Semi-Markov Model of Symbolic Music for Score Following," in *Proc. of the 16th Int. Society for Music Information Retrieval Conference (ISMIR)*, 2015.

- [10] D. Schwarz, N. Orio, and N. Schnell, "Robust Polyphonic Midi Score Following with Hidden Markov Models," in *Int. Computer Music Conference (ICMC)*, 2004.
- [11] S. Dixon, "An On-Line Time Warping Algorithm for Tracking Musical Performances," in Proc. of the 19th Int. Joint Conference on Artificial Intelligence (IJCAI), 2005.
- [12] M. Dorfer, A. Arzt, and G. Widmer, "Towards Score Following in Sheet Music Images," in Proc. of the 17th Int. Society for Music Information Retrieval Conference (ISMIR), 2016.
- [13] M. Dorfer, F. Henkel, and G. Widmer, "Learning to Listen, Read, and Follow: Score following as a Reinforcement Learning Game," in *Proc.* of the 19th Int. Society for Music Information Retrieval Conference (ISMIR), 2018.
- [14] F. Henkel, S. Balke, M. Dorfer, and G. Widmer, "Score Following as a Multi-Modal Reinforcement Learning Problem," *Trans. of the Int. Society for Music Information Retrieval*, vol. 2, no. 1, 2019.
- [15] F. Henkel, R. Kelz, and G. Widmer, "Learning to Read and Follow Music in Complete Score Sheet Images," in *Proc. of the 21st Int. Society for Music Information Retrieval Conference (ISMIR)*, 2020.
- [16] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] M. Dorfer, J. Hajič jr., A. Arzt, H. Frostel, and G. Widmer, "Learning Audio–Sheet Music Correspondences for Cross-Modal Retrieval and Piece Identification," *Trans. of the Int. Society for Music Information Retrieval*, vol. 1, no. 1, 2018.
- [18] F. Henkel, R. Kelz, and G. Widmer, "Audio-Conditioned U-Net for Position Estimation in Full Sheet Images," in *Proc. of the 2nd Int. Workshop on Reading Music Systems*, 2019.
- [19] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "FiLM: Visual reasoning with a general conditioning layer," in *Proc. of the 32nd* AAAI Conference on Artificial Intelligence, 2018.
- [20] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in Advances in Neural Information Processing Systems (NeurIPS), 2015.
- [21] R. Hu, M. Rohrbach, and T. Darrell, "Segmentation from natural language expressions," in *Proc. of the 14th Eur. Conference on Computer Vision (ECCV)*, 2016.
- [22] T. Ridnik, H. Lawen, A. Noy, and I. Friedman, "TResNet: High Performance GPU-Dedicated Architecture," arXiv preprint arXiv:2003.13630, 2020.
- [23] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer Normalization," arXiv preprint arXiv:1607.06450, 2016.
- [24] D. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)," in *Proc. of the* 4th Int. Conference on Learning Representations (ICLR), 2016.
- [25] J. Redmon and A. Farhadi, "YOLO9000: Better, Faster, Stronger," in Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- [26] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in Proc. of the 32nd Int. Conference on Machine Learning (ICML), 2015.
- [27] T. Grill and J. Schlüter, "Two convolutional neural networks for bird detection in audio signals," in *Proc. of the 25th Eur. Signal Processing Conference (EUSIPCO)*, 2017.
- [28] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, 1997.
- [29] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," in Proc. of the 7th Int. Conference on Learning Representations (ICLR), 2019.
- [30] T. He, Z. Zhang, H. Zhang, Z. Zhang, J. Xie, and M. Li, "Bag of Tricks for Image Classification with Convolutional Neural Networks," in *Proc.* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
- [31] A. M. Saxe, J. L. McClelland, and S. Ganguli, "Exact solutions to the nonlinear dynamics of learning in deep linear neural networks," *arXiv* preprint arXiv:1312.6120, 2013.
- [32] F. Gers, J. Schmidhuber, and F. Cummins, "Learning to Forget: Continual Prediction with LSTM," *Neural Computation*, vol. 12, no. 10, 2000.
- [33] I. Loshchilov and F. Hutter, "SGDR: Stochastic Gradient Descent with Warm Restarts," in Proc. of the 5th Int. Conference on Learning Representations (ICLR), 2017.
- [34] M. Müller, Fundamentals of Music Processing. Springer Verlag, 2015.