

# Reconstruction Meets Prediction: Complementary SSL Objectives for Music Transcription

Mary Pilataki, Matthias Mauch, and Simon Dixon

**Abstract**—Multi-instrument automatic music transcription (MIAMT) requires learning representations from limited labelled data. We present a controlled comparison of two self-supervised paradigms: masked spectrogram reconstruction with a masked autoencoder (MAE) and predictive learning with a joint embedding predictive architecture (JEPA). We use identical Transformer encoders, evaluating both transcription performance and latent space properties: geometric organisation, retrieval quality and representational capacity. On transcription, the two paradigms are broadly comparable, with task-specific advantages: JEPA performs better for monophonic wind instruments, MAE for dense polyphonic ensembles. Both substantially exceed the prior SSL MIAMT baseline, and on URMP both also exceed a recent state-of-the-art supervised system. In latent space, JEPA achieves better geometric organisation while MAE preserves higher representational capacity. Furthermore, MAE converges substantially faster under supervision. Initialising JEPA from MAE-pretrained weights consistently improves transcription, demonstrating that the two objectives are complementary rather than competing. Three key findings emerge. First, no single latent metric reliably predicts transcription performance. Second, MAE provides a richer initialisation for predictive pretraining than random. Third, reconstruction and prediction objectives reveal opposite capacity trajectories under supervision, a behaviour that persists regardless of initialisation. Code and pretrained models will be released on GitHub.<sup>1</sup>

**Index Terms**—music transcription, self-supervised learning, masked autoencoders, joint embedding predictive architectures

## I. INTRODUCTION

Multi-instrument automatic music transcription (MIAMT) extracts note-level symbolic representations from polyphonic audio containing multiple instruments. It is fundamental to Music Information Retrieval (MIR), with applications in music education, interactive music systems and musicology. Despite deep learning advances, accurate transcription remains challenging due to overlapping harmonics, timbral diversity of source instruments and interrelated subtasks, including multi-pitch estimation (MPE), onset/offset detection and instrument recognition (IR) [1].

State-of-the-art systems rely on supervised learning with spectral features [2], [3], [4], [5], [6] but face a critical bottleneck: the scarcity of large-scale, high-quality annotated datasets. Manual note-level annotation requires expert musicians and is expensive, resulting in limited instrument and

genre coverage. This constrains supervised approaches that depend on large labeled corpora.

Self-supervised learning (SSL) offers a promising alternative by learning representations from unlabeled data, with masked modeling being a dominant paradigm in the audio domain [7], [8], [9], [10], [11], [12]. Generative approaches like the Audio Masked Autoencoder (AudioMAE) [10], [13] and hybrid strategies that combine complementary discriminative and multimodal objectives [12], [14], [15], have shown strong downstream performance. Predictive frameworks based on Joint Embedding Predictive Architectures (JEPA), including Audio-JEPA [9] and music-specific variants [16], [11] and structured methods like MATPAC [17], [18] have shifted the focus towards learning directly in the representation space, bypassing signal reconstruction. However, how these self-supervised objectives shape the learned latent space and whether the resulting geometric and capacity differences translate into measurable performance differences remains unclear.

We hypothesise that the SSL objective fundamentally shapes both representation geometry and representational capacity, and that analysing these properties can reveal complementary strengths that transcription scores alone cannot expose. To test this, we present a controlled comparison of two SSL frameworks for MIAMT based on JEPA [8], [9], [11] and AudioMAE [10]. Both use identical Transformer-based spectrogram encoders and published hyperparameters, differing only in masking strategy and training objective. JEPA models high-level semantic relationships between contextual and target blocks directly in latent space, without signal reconstruction [8], while AudioMAE reconstructs missing spectrogram patches from the surrounding context, encouraging detailed spectral modelling [10], [19]. This controlled setup allows us to attribute any differences in latent structure and downstream performance to the SSL objective itself.

While MIR studies typically evaluate SSL models mainly through downstream performance alone, such *black box* evaluation provides limited insight into representation structure. Inspired by recent encoder assessment frameworks [20], [21] and representation analyses used in vision [22] and audio [23] SSL, we extend evaluation beyond transcription metrics. Specifically, we measure geometric organisation via alignment and uniformity of the hypersphere [24], pitch-space neighbourhood structure via pitch locality, retrieval quality via kNN accuracy [22] and representational capacity by the RankMe score [21]. This analysis reveals how musical information is structured and utilised in the latent space, rather than relying on downstream scores alone.

Our main findings are as follows. Transcription performance

Manuscript submitted for review on 19 March 2026.

Simon Dixon and Mary Pilataki are with the Centre for Digital Music, Queen Mary University of London, UK (e-mail: s.e.dixon@qmul.ac.uk; m.pilataki-manika@qmul.ac.uk).

Matthias Mauch is with Apple, UK (email: mmauch@apple.com).

<sup>1</sup><https://github.com/mariyapilataki/amt-mae>




Fig. 1. SSL MIAMT framework: MIAMT-MAE pretraining (top), MIAMT-JEPA pretraining (middle) and the shared finetuning setup (bottom).

is broadly comparable between the two models, with task-specific advantages. MIAMT-JEPA outperforms on monophonic wind instruments while MIAMT-MAE is stronger on dense polyphonic content (piano and ensemble). Both substantially exceed the prior SSL transcription baseline [11], and on URMP both also exceed a recent state-of-the-art supervised system. MAE-pretrained initialisation of JEPA (MIAMT-JEPA<sup>†</sup>) yields consistent gains across all tasks, demonstrating that the two objectives are complementary and that MAE provides a powerful initialisation for predictive pretraining. Additionally, MAE converges substantially faster than JEPA under supervised finetuning, an important practical consideration. Latent-space analysis reveals complementary trade-offs. MIAMT-JEPA achieves superior geometric organisation (uniformity, pitch locality) and retrieval quality, while MIAMT-MAE preserves higher representational capacity. Surprisingly, finetuning produces opposite capacity trajectories, with MAE expanding and JEPA reducing, a behaviour that persists regardless of initialisation. No single latent metric reliably predicts transcription performance, indicating that the two SSL objectives are best understood as complementary rather than competing.

## II. MODEL ARCHITECTURES

Both MIAMT-MAE and MIAMT-JEPA use identical Vision Transformer (ViT-B) encoders [25] adapted for spectrograms, differing only in pretraining modules (Figure 1). Input log-mel spectrograms are divided into non-overlapping  $[8 \times 8]$  patches and are linearly projected to a 768-dimensional embedding space with learnable positional encodings that retain

temporal and spectral order. The sequence is processed by 12 Transformer blocks, each with 12 self-attention heads, a feed-forward layer with expansion ratio 4 and Layer Normalisation. The normalised output hidden states form the encoder representation, which is input to the transcription heads.

During pretraining, MIAMT-MAE uses an 8-block decoder (512-dim) to reconstruct masked patches [10], while MIAMT-JEPA uses a 6-block predictor (384-dim) to estimate target embeddings from an Exponential Moving Average (EMA) encoder in latent space [8]. Both pretraining modules are discarded at finetuning.

Following the MIAMT-JEPA baseline [11], both models share a shallow transcription head consisting of two hidden linear layers with sigmoid activations, producing pitch activity probabilities per instrument. The head accepts 128-dimensional inputs (one per mel bin) and produces an  $[88 \times 11]$  matrix covering 88 pitches (A0-C8) and 11 instrument classes. Latent embeddings are projected into a spectrogram space via transposed convolution before being passed to the transcription head (Figure 1, bottom).

## III. EXPERIMENTAL SETUP

### A. Datasets and Input Features

Following [11], SSL pretraining uses GTZAN [26], the Violin Bach dataset [27], Mazurkas<sup>2</sup>, and the FrançoisLeduc guitar dataset [28], totalling approximately 159 hours of unlabeled audio. Supervised finetuning uses Slakh [29], MusicNet’s EM version [30], [31], GuitarSet [32], URMP [33] and MAESTRO [34] with standard train-test splits from [3], [2].

Log-mel spectrograms are computed from 44.1 kHz mono audio using 128 mel bands, a 100 ms Hanning window, 10 ms hop size and a 20 Hz lower frequency bound, using torchaudio’s *MelSpectrogram* transform [35]. This yields  $[128 \times 128]$  representations corresponding to 1.28 s audio segments. Consecutive segments overlap by 25% (0.32 s stride). Spectrograms are normalised using dataset statistics [10].

### B. Experiments

Both models are pretrained self-supervised for 300 epochs using their published configurations [10], [8], [9]. We then finetune each encoder jointly with the multi-instrument transcription head, discarding SSL-specific modules (decoder, predictor) and objectives and optimising only the transcription loss. No encoder layers are frozen during finetuning. Preliminary experiments showed that this transcription-loss-only finetuning setup outperforms frozen-encoder probing. In [11], three variants were compared, end-to-end training, joint-loss finetuning, and frozen-encoder probing, with the frozen probe best. We attribute this reversal to dropping the SSL objective during finetuning. MIAMT-MAE converges in 100 epochs, while MIAMT-JEPA requires 300. This  $3 \times$  difference in finetuning time suggests MAE representations align more quickly with the transcription objective. We report results from the best validation checkpoint for each model. We additionally evaluate warm-start initialisation by pretraining

<sup>2</sup><http://www.charm.rhul.ac.uk/index.html>

the JEPA objective from MAE pretrained encoder weights, denoting this model MIAMT-JEPA<sup>†</sup>. All hyperparameters are kept identical to MIAMT-JEPA.

The reverse configuration (JEPA encoder weights initialising MAE pretraining) yields representational collapse ( $\mathcal{R}=79$ ), a sharp drop in effective dimensionality, and is not investigated further. This failure likely stems from the compressed, semantically invariant nature of JEPA representations ( $\mathcal{R}=224$ ), lacking the fine-grained spectral detail necessary to initialise a reconstruction-based decoder, leading to immediate representational collapse.

### C. Objectives and Training Details

MIAMT-JEPA minimises smooth  $\ell_1$  (Huber) loss between predictor outputs and EMA target representations [8], [11], [9]. MIAMT-MAE optimises Mean Squared Error (MSE) reconstruction loss over masked spectrogram patches [10]. These objectives reflect different inductive biases. MIAMT-MAE focuses on pixel-space fidelity, encouraging the encoder to capture fine-grained spectral detail. MIAMT-JEPA operates entirely in representation space, bypassing signal reconstruction in favour of semantic consistency between context and target embeddings.

Both models are finetuned with an instrument-aware transcription objective [11] combining pitch activity and onset detection:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{trans}} + \alpha \mathcal{L}_{\text{onset}}, \quad \alpha = 0.5. \quad (1)$$

$\mathcal{L}_{\text{trans}}$  applies weighted binary cross-entropy over the pitch activity matrix  $\mathbf{Y} \in \{0,1\}^{T \times N \times I}$ , with instrument-activity weights  $w_{t,i}^{\text{inst}} \in \{0.70, 0.30\}$  (0.70 when instrument  $i$  active, 0.30 otherwise).  $T$ ,  $N$ , and  $I$  denote the number of frames, pitches, and instrument classes respectively:

$$\mathcal{L}_{\text{trans}} = \frac{1}{TNI} \sum_{t,n,i} w_{t,i}^{\text{inst}} \ell_{\text{BCE}}(\hat{y}_{t,n,i}, y_{t,n,i}). \quad (2)$$

$\mathcal{L}_{\text{onset}}$  upweights note boundaries with a  $\pm 1$ -frame ( $\pm 10$  ms) tolerance window around exact onsets to allow for minor annotation timing imprecision. Weights are  $w^{\text{on}}=0.90$  (onset) and  $w^{\text{off}}=0.10$  (non-onset) [6]:

$$\mathcal{L}_{\text{onset}} = \frac{\sum_{t,n,i} w_{t,n,i}^{\text{on/off}} \ell_{\text{BCE}}(\hat{y}_{t,n,i}, y_{t,n,i})}{\sum_{t,n,i} w_{t,n,i}^{\text{on/off}} + \epsilon}, \quad (3)$$

where  $\epsilon=10^{-6}$  prevents division by zero.

### D. Latent Space Evaluation

Latent representations are extracted by mean-pooling encoder outputs across the time axis followed by  $\ell_2$  normalisation. The NSynth test set [36] is used for evaluation, as it provides controlled monophonic recordings with ground-truth pitch and instrument labels. We report four metric categories. **(1) Representational capacity** quantifies how many effective dimensions are used. RankMe  $\mathcal{R}$  [21] computes the exponential of Shannon entropy over normalised singular values (range: 1 to embedding dimension; higher indicates more dimensions actively used, less representational collapse).

**(2) Geometric organisation** measures hypersphere structure via alignment  $\mathcal{A}$  and uniformity  $\mathcal{U}$  [24]. Alignment computes expected  $\ell_2$  distance between same-class pairs (range: 0-2; lower means tighter within-class clusters). Uniformity computes log-average pairwise Gaussian potential (typically  $-4$  to 0; more negative means more evenly distributed across the hypersphere, avoiding collapsed representations). Both metrics operate on  $\ell_2$ -normalised embeddings on the unit hypersphere. **(3) Pitch locality**  $L_p(\Delta)$  measures, for each query embedding, the fraction of  $k$ -nearest neighbours in embedding space whose ground truth pitch falls within  $\Delta$  semitones of the query, averaged over all queries (range: 0-1; higher means same-pitch excerpts cluster nearby). We report results for  $\Delta=1$  (semitone) and  $\Delta=12$  (octave). **(4) Retrieval quality** is measured by weighted k-NN accuracy  $\kappa$  [22] for pitch and instrument classification using temperature-scaled cosine similarity voting ( $k=10$ ,  $T=0.07$ ). Higher accuracy indicates embeddings that support better nearest-neighbour-based predictions.

### E. Transcription Evaluation

Multi-instrument note-wise F1-scores are computed using `mir_eval` [37]. Frame-level predictions are post-processed using the method by Bittner et al. [6] with a threshold of 0.3, removing events shorter than 130 ms. Notes match if pitch is within 0.5 semitones, onset within 50 ms and instrument prediction is correct [2], [3]. We compare with MT3 [2], YourMT3 [3] and the MIAMT-JEPA baseline [11]. Direct comparisons should be interpreted cautiously due to differences in training data, model scale and instrument vocabularies across systems.

## IV. RESULTS AND DISCUSSION

### A. Learned Embeddings

Table I compares how reconstruction-based (MIAMT-MAE) and predictive (MIAMT-JEPA) objectives structure the latent space in pretrained (PT) and finetuned (FT) states. The two models exhibit complementary trade-offs across geometric organisation, retrieval quality and representational capacity.

MIAMT-JEPA achieves superior uniformity ( $\mathcal{U}=-0.610$  vs.  $-0.372$ ) and pitch locality ( $L_p(1)=0.371$  vs.  $0.240$ ), indicating well-spread embeddings with better pitch-space organisation. MIAMT-MAE achieves tighter within-class alignment ( $\mathcal{A}_p=0.195$  vs.  $0.313$ ), meaning same-class samples are closer together despite less uniform distribution. MIAMT-JEPA substantially outperforms on k-NN accuracy for both pitch ( $\kappa_p=71.9\%$  vs.  $58.3\%$ ,  $+13.6\%$ ) and instrument ( $\kappa_i=52.9\%$  vs.  $47.9\%$ ,  $+5.0\%$ ), confirming that superior uniformity and pitch locality translate directly to better retrieval.

MIAMT-MAE achieves higher effective dimensionality ( $\mathcal{R}=348$  vs.  $200$ ). This indicates that embeddings span a broader portion of the latent space. Surprisingly, finetuning causes the two models to diverge. MIAMT-MAE expands ( $\mathcal{R}: 315 \rightarrow 348$ ) while MIAMT-JEPA reduces ( $\mathcal{R}: 224 \rightarrow 200$ ), suggesting different adaptation mechanisms. MIAMT-MAE begins with task-relevant structure and undergoes minimal geometric change under supervision (all  $|\Delta| < 3\%$  except capacity  $+10\%$ ). MIAMT-JEPA begins with loose alignment

TABLE I

LATENT SPACE EVALUATION ON NSYNTH FOR PRETRAINED (PT) AND FINETUNED (FT) ENCODERS. GEOMETRY IS MEASURED VIA ALIGNMENT  $\mathcal{A}$ , UNIFORMITY  $\mathcal{U}$ , AND PITCH LOCALITY  $L_p(\Delta)$ . RETRIEVAL QUALITY VIA K-NN ACCURACY  $\kappa$  AND REPRESENTATIONAL CAPACITY VIA RANKME  $\mathcal{R}$ .  $\dagger$ INITIALISED FROM MAE PRETRAINED ENCODER WEIGHTS PRIOR TO JEPa PRETRAINING.

| Model                 | Ckpt | Geometry                   |                          |                   |                    | k-NN Acc. (%)       |                     | Capacity               |
|-----------------------|------|----------------------------|--------------------------|-------------------|--------------------|---------------------|---------------------|------------------------|
|                       |      | $\mathcal{A}_p \downarrow$ | $\mathcal{U} \downarrow$ | $L_p(1) \uparrow$ | $L_p(12) \uparrow$ | $\kappa_p \uparrow$ | $\kappa_i \uparrow$ | $\mathcal{R} \uparrow$ |
| MIAMT-MAE             | PT   | 0.197                      | -0.360                   | 0.243             | 0.711              | 60.7                | 48.8                | 315                    |
|                       | FT   | 0.195                      | -0.372                   | 0.240             | 0.704              | 58.3                | 47.9                | 348                    |
| MIAMT-JEPa            | PT   | 0.541                      | -1.119                   | 0.307             | 0.693              | 64.8                | 48.7                | 224                    |
|                       | FT   | 0.313                      | -0.610                   | 0.371             | 0.753              | 71.9                | 52.9                | 200                    |
| MIAMT-JEPa $^\dagger$ | PT   | 0.137                      | -0.250                   | 0.227             | 0.724              | 56.3                | 51.1                | 282                    |
|                       | FT   | 0.189                      | -0.312                   | 0.227             | 0.720              | 58.3                | 49.0                | 212                    |

TABLE II

TRANSCRIPTION F1-SCORES. COLUMNS GROUPED BY DATASET AND INSTRUMENT.

| Model                 | <i>Maestro</i> | <i>GuitarSet</i> | <i>MusicNet</i> |             | <i>URMP</i> |
|-----------------------|----------------|------------------|-----------------|-------------|-------------|
|                       | Piano          | Guitar           | Strings         | Winds       | Ensemble    |
| MIAMT-MAE             | 82.9           | 82.1             | 69.6            | 56.7        | 74.2        |
| MIAMT-JEPa            | 81.1           | 82.4             | 69.8            | 60.3        | 72.0        |
| MIAMT-JEPa $^\dagger$ | 85.3           | 84.3             | 74.7            | 65.3        | 78.2        |
| SSL baseline          | -              | 63.5             | 47.9            | 35.7        | <b>78.9</b> |
| YourMT3               | <b>97.0</b>    | <b>91.7</b>      | <b>91.3</b>     | <b>83.5</b> | 68.0        |
| MT3                   | 94.9           | 78.0             | -               | -           | -           |

( $\mathcal{A}_p=0.541$ ) and high uniformity ( $\mathcal{U}=-1.119$ ). During fine-tuning it undergoes dramatic reorganisation. Alignment tightens  $-42\%$ , uniformity increases  $+45\%$  and k-NN accuracy improves  $+7.1\%$ . However, capacity reduces  $-11\%$ , indicating that reconstruction pretraining produces immediately usable representations while prediction yields initially task-misaligned but adaptable ones.

MAE-pretrained initialisation of JEPa (MIAMT-JEPa $^\dagger$ ) produces intermediate pretrained capacity ( $\mathcal{R}=282$ ) and tighter initial alignment ( $\mathcal{A}_p=0.137$  vs.  $0.541$  compared to scratch JEPa), suggesting MAE initialisation preserves spectral structure that constrains early representation geometry. Finetuning reduces capacity ( $282 \rightarrow 212$ ), closely mirroring scratch JEPa’s trajectory ( $224 \rightarrow 200$ ). This indicates that the JEPa objective governs adaptation dynamics regardless of initialisation.

MIAMT-JEPa performs better on 5 out of 7 latent metrics while MIAMT-MAE achieves nearly double the representational capacity. However, as shown in Section IV-B, neither property alone reliably predicts transcription performance.

### B. Transcription Performance

Table II reports multi-instrument transcription F1-scores. Both MIAMT-MAE and MIAMT-JEPa substantially outperform the SSL transcription baseline [11] on three of the datasets. Despite using only 159 hours of unlabeled pretraining data, a base ViT architecture, and no dataset-specific augmentation, both models also exceed YourMT3 [3] on the URMP data (74.2% and 72.0% vs. 68.0%). Transcription results are broadly comparable between the two scratch models, with

neither consistently dominating. MIAMT-JEPa outperforms on wind instruments (MusicNet Winds: 60.3 vs. 56.7,  $+3.6\%$ ) and narrowly on strings and guitar. MIAMT-MAE is stronger on piano (Maestro: 82.9 vs. 81.1,  $+1.8\%$ ) and ensemble transcription (URMP: 74.2 vs. 72.0,  $2.2\%$ ). These task-specific differences suggest complementary strengths rather than a clear winner.

MIAMT-JEPa $^\dagger$  outperforms both scratch models across all tasks, with gains of 1.9-8.6% F1 over the stronger scratch baseline per task. MIAMT-JEPa $^\dagger$  achieves higher piano (85.3 vs. 82.9) and ensemble (78.2 vs. 74.2) scores than MIAMT-MAE despite lower capacity ( $\mathcal{R}=212$  vs. 348). This challenges the interpretation that representational capacity predicts downstream performance [21]. MAE’s spectrogram-grounded structure appears directly exploitable by the JEPa objective, producing features neither objective achieves alone. These results indicate that no single latent metric reliably predicts transcription performance, and that the two SSL objectives are better understood as complementary than competing.

To verify that transcription gains stem from SSL pretraining, we additionally train encoders jointly with the transcriber from scratch, using only the transcription loss. The performance gap is largest on URMP (9.8% average decline), the smallest labelled corpus and a multi-instrument setting, and smallest on Maestro (3.2% average decline), a large single-instrument piano dataset. This indicates that SSL pretraining is most beneficial when labelled data is scarce and the task is more complex while its benefit diminishes as labelled data grows. These findings are consistent with prior work reporting that SSL is most valuable in low-resource regimes [38], [39], and that supervised-only training can approach SSL performance given sufficient labelled data [40].

## V. CONCLUSION

We compared reconstruction-based (MIAMT-MAE) and predictive (MIAMT-JEPa) SSL for multi-instrument transcription. Both outperform the prior SSL baseline and exceed a state-of-the-art supervised system on URMP. Supervised-only ablations confirm these gains stem from SSL pretraining rather than the encoder architecture alone, with the benefit largest where labelled data is scarce. The SSL objectives prove complementary. JEPa is better on monophonic wind instruments, MAE on dense polyphonic content and initialising JEPa pretraining from MAE-pretrained weights (MIAMT-JEPa $^\dagger$ ) results in further consistent gains. Three insights emerge. First, no single latent metric reliably predicts transcription performance. RankMe explains differences among from-scratch models but not cross-initialisation gains. Second, MAE-pretrained initialisation provides structure directly useful for predictive pretraining. Third, the objectives exhibit opposite capacity trajectories under supervision, regardless of initialisation. Unlike the vision domain, the fact that JEPa does not dominate here suggests optimal SSL strategies are strongly modality- and task-dependent. Constrained to ViT-Base with fixed masking, our results motivate hybrid objectives and initialisation strategies that combine MAE’s capacity preservation with JEPa’s geometric organisation.

## REFERENCES

- [1] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic music transcription: An overview," *IEEE Signal Processing Magazine*, vol. 36, no. 1, pp. 20–30, 2019.
- [2] J. Gardner, I. Simon, E. Manilow, C. Hawthorne, and J. Engel, "MT3: Multi-task multitrack music transcription," in *Proc. ICLR*, 2021.
- [3] S. Chang, E. Benetos, H. Kirchhoff, and S. Dixon, "YourMT3+: Multi-instrument music transcription with enhanced transformer architectures and cross-dataset stem augmentation," in *Proc. MLSP*, 2024.
- [4] C. Hawthorne, E. Elsen, J. Song, A. Roberts, I. Simon, C. Raffel, J. Engel, S. Oore, and D. Eck, "Onsets and frames: Dual-objective piano transcription," in *Proc. ISMIR*, 2018.
- [5] Q. Kong, B. Li, X. Song, Y. Wan, and Y. Wang, "High-resolution piano transcription with pedals by regressing onset and offset times," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2020.
- [6] R. M. Bittner, J. J. Bosch, D. Rubinstein, G. Meseguer-Brocal, and S. Ewert, "A lightweight instrument-agnostic model for polyphonic note transcription and multipitch estimation," in *Proc. ICASSP*, 2022.
- [7] C. Doersch, A. Gupta, and A. A. Efros, "Unsupervised visual representation learning by context prediction," in *Proc. ICCV*, 2015, pp. 1422–1430.
- [8] M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas, "Self-supervised learning from images with a joint-embedding predictive architecture," in *Proc. CVPR*, 2023.
- [9] L. Tuncay, E. Labbé, E. Benetos, and T. Pellegrini, "Audio-JEPA: Joint-embedding predictive architecture for audio representation learning," in *Proc. ICME*, 2025.
- [10] P.-Y. Huang, H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, and C. Feichtenhofer, "Masked autoencoders that listen," in *Proc. NeurIPS*, 2022.
- [11] M. Pilataki, M. Mauch, and S. Dixon, "Self-supervised learning with a JEPA framework for multi-instrument music transcription," in *Proc. WASPAA*, 2025.
- [12] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, "Masked modeling duo: Towards a universal audio pre-training framework," *IEEE Trans. Audio, Speech, Lang. Process.*, 2024.
- [13] S. Yadav, S. Theodoridis, and Z.-H. Tan, "Audiomae++: Learning better masked audio representations with swiglu ffns," in *Proc. MLSP*, 2025.
- [14] D. Niizumi, D. Takeuchi, Y. Ohishi, N. Harada, and K. Kashino, "M2d-clap: Masked modeling duo meets clap for learning general-purpose audio-language representation," in *Proc. Interspeech*, 2024.
- [15] Y.-C. Huang, Y. Gong, H. Jiang, C.-K. Chuang, A. Shrivastava, and J. Glass, "Contrastive audio-visual masked autoencoder," in *Proc. ICLR*, 2023.
- [16] A. Riou, S. Lattner, G. Hadjeres, M. Anslow, and G. Peeters, "Stem-JEPA: A joint-embedding predictive architecture for musical stem compatibility estimation," in *Proc. ISMIR*, 2024.
- [17] A. Quelenec, P. Chouteau, G. Peeters, and S. Essid, "Masked latent prediction and classification for self-supervised audio representation learning," in *Proc. ICASSP*, 2025.
- [18] —, "MATPAC++: Enhanced masked latent prediction for self-supervised audio representation learning," *arXiv preprint arXiv:2508.12709*, 2025.
- [19] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. CVPR*, 2022.
- [20] J. Zhang, H. Dinkel, Y. Niu, C. Liu, S. Cheng, A. Zhao, and J. Luan, "X-ARES: A comprehensive framework for assessing audio encoder performance," in *Proc. Interspeech*, 2025.
- [21] Q. Garrido, R. Balestrierio, L. Najman, and Y. Lecun, "Rankme: Assessing the downstream performance of pretrained self-supervised representations by their rank," in *Proc. ML*, 2023.
- [22] M. Caron, H. Touvron, I. Misra, H. Jegou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proc. ICCV*, 2021.
- [23] H. Dinkel, Z. Yan, Y. Wang, J. Zhang, Y. Wang, and B. Wang, "Scaling up masked audio encoder learning for general audio classification," in *Proc. Interspeech*, 2024.
- [24] T. Wang and P. Isola, "Understanding contrastive representation learning through alignment and uniformity on the hypersphere," in *Proc. ICML*, 2020.
- [25] A. Kolesnikov, A. Dosovitskiy, D. Weissenborn, G. Heigold, J. Uszkoreit, L. Beyer, M. Minderer, M. Dehghani, N. Houlsby, S. Gelly, T. Unterthiner, and X. Zhai, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2021.
- [26] G. Tzanetakis and P. Cook, "Musical genre classification of audio signals," *IEEE Trans. Audio, Speech, Lang. Process.*, 2002.
- [27] H.-W. Dong, C. Zhou, T. Berg-Kirkpatrick, and J. McAuley, "Deep performer: Score-to-audio music performance synthesis," in *Proc. ICASSP*, 2022.
- [28] X. Riley, D. Edwards, and S. Dixon, "High resolution guitar transcription via domain adaptation," in *Proc. ICASSP*, 2024.
- [29] E. Manilow, G. Wichern, P. Seetharaman, and J. Le Roux, "Cutting music source separation some Slakh: A dataset to study the impact of training data quality and quantity," *Proc. WASPAA*, 2019. [Online]. Available: <http://dx.doi.org/10.1109/WASPAA.2019.8937170>
- [30] J. Thickstun, Z. Harchaoui, and S. Kakade, "Learning features of music from scratch," in *Proc. ICLR*, 2017.
- [31] B. Maman and A. H. Bermano, "Unaligned supervision for automatic music transcription in the wild," in *Proc. ICML*, 2022.
- [32] Q. Xi, R. M. Bittner, J. Pauwels, X. Ye, and J. P. Bello, "Guitarset: A dataset for guitar transcription," in *Proc. ISMIR*, 2018.
- [33] B. Li, X. Liu, K. Dinesh, Z. Duan, and G. Sharma, "Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications," *IEEE Trans. Multimedia*, vol. 21, no. 2, pp. 522–535, 2019.
- [34] C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C.-Z. A. Huang, S. Dieleman, E. Elsen, J. Engel, and D. Eck, "Enabling factorized piano music modeling and generation with the MAESTRO dataset," in *Proc. ICLR*, 2019.
- [35] J. Hwang, M. Hira, C. Chen, X. Zhang, Z. Ni, G. Sun, P. Ma, R. Huang, V. Pratap, Y. Zhang, A. Kumar, C.-Y. Yu, C. Zhu, C. Liu, J. Kahn, M. Ravanelli, P. Sun, S. Watanabe, Y. Shi, and Y. Tao, "Torchaudio 2.1: Advancing speech recognition, self-supervised learning, and audio processing components for pytorch," in *Proc. ASRU*, 2023.
- [36] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with wavenet autoencoders," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Eds. PMLR, 2017.
- [37] C. Raffel, B. McFee, E. J. Humphrey, J. Salamon, O. Nieto, D. Liang, and D. P. W. Ellis, "mir eval: A transparent implementation of common MIR metrics," in *Proc. ISMIR*, 2014.
- [38] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. NeurIPS*, 2020.
- [39] A. Newell and J. Deng, "How useful is self-supervised pretraining for visual tasks?" in *Proc. CVPR*, 2020.
- [40] K. He, R. Girshick, and P. Dollár, "Rethinking imagenet pre-training," in *Proc. ICCV*.