





Analysis of Coordination and Multi-Agent Dynamics in Sequential and Generative Models

Shlomo Dubnov






European Research Council
Established by the European Commission



Snapshots from our research:

Audio Representation Learning → Downstream Applications



Lower Down Barriers for Human to Understand and Create Audio and Music

HTS-AT: A Hierarchical Token-Semantic Audio Transformer for Sound Classification and Detection

Publisher: **IEEE**

Cite This

PDF


[Ke Chen](#) ; [Xingjian Du](#) ; [Bilei Zhu](#) ; [Zejun Ma](#) ; [Taylor Berg-Kirkpatrick](#) ; [Shlomo Dubnov](#) All Authors

283

Cites in
Papers

3492

Full
Text Views



[1] Ke Chen et al. "HTS-AT: A Hierarchical Token-Semantic Audio Transformer for Sound Classification and Detection", ICASSP 2022.



Conferences > ICASSP 2023 - 2023 IEEE Inter... ?

Large-Scale Contrastive Language-Audio Pretraining with Feature Fusion and Keyword-to-Caption Augmentation

Publisher: **IEEE**

[Cite This](#)



PDF

[Yusong Wu](#) ; [Ke Chen](#) ; [Tianyu Zhang](#) ; [Yuchen Hui](#) ; [Taylor Berg-Kirkpatrick](#) ; [Shlomo Dubnov](#) All Authors

440

Cites in
Papers

4537

Full
Text Views



Mila, Québec Artificial Intelligence Institute, Université de Montréal
University of California San Diego
LAION

MusicLDM: Enhancing Novelty in text-to-music Generation Using Beat-Synchronous mixup Strategies

Publisher: IEEE

[Cite This](#)

[PDF](#)


[Ke Chen](#) ; [Yusong Wu](#) ; [Haohe Liu](#) ; [Marianna Nezhurina](#) ; [Taylor Berg-Kirkpatrick](#) ; [Shlomo Dubnov](#) All Authors

73

Cites in
Papers

1346

Full
Text Views




<https://musicldm.github.io/>

The OMax Improvisation Paradigm



Presented by the OMax Brothers
Assayag, Bloch, Chemillier, Dubnov



Talk Plan

- Why Robustness is tied to Information?
- Directed Information and Information Rate
- Music Information Dynamics
 - Examples: Texture and Counterpoint
- DI Estimation using Foundation Music models
- Information in Control and Perception-Action
- Robustness in Control: Failure conditions and animated replay
- Quadruped Example
- Conclusions

Why Robustness is tied to Information?

Robustness is the ability to identify and preserve functionally useful information structure under changes in data, environment, representation, or task.

Uncertainty→Prediction→Interaction→Function


$H(X) \rightarrow H(X_t|X_{<t}) \rightarrow I(X \rightarrow Y) \rightarrow$ information needed/used for action

Robustness = preserving useful structure under change

The goal is not invariance of the data—it is invariance of the information structure needed for prediction, coordination, or function.

Robust AI: What Information Must Survive?

A perception–action view: function should survive changes in observations, environment, or components



Generative modeling

learn $p_{\theta}(X)$ Fit and sample the observed data distribution

Robust-AI question

Which information remains useful for prediction and action when the conditions change?

Robustness: preserve functionally relevant information — not the exact observations

Prediction, Compression, and Control

Three information-theoretic views of the same perception–action loop

Relevant Information / Bottleneck

What must I know?

$$\min_{\pi} [I(S; A) - \beta \mathbb{E} U]$$

Minimum state information reflected in the action for a desired level of performance.

Predictive Information

What can I predict?

$$I(S_{\text{past}}; S_{\text{future}})$$

How much does the sensory past tell us about the sensory future?

A task-independent view of predictive structure.

Empowerment

What can I control?

$$\mathcal{E}_T(s) = \max_{p(a_{0:T-1})} I(A_{0:T-1}; S_T | S_0 = s)$$


Potential influence of possible action sequences on future observable states.

A measure of options / control capacity.

What is relevant? • What is predictable? • What is controllable?


From One Agent to Interacting Components

Generalize the perception–action view to coordination, multi-agent dynamics, and feedback channels




Information Rate

- Entropy & Information



- Communication Channel



$$I(x, y) = H(x) - H(x | y)$$

y – past experience, what you heard so far

x – new material


$H(x)$ – uncertainty about x

$H(x|y)$ – uncertainty about x when we know already y

*If y is long enough, $H(x|y)$ becomes **Entropy Rate***

$I(x,y)$ – how much the past tells us about the future

Music Information Dynamics




Validation and violation of expectations is a technique commonly used by composers

$$\rho(x) = I(x_{\text{present}}, x_{\text{past}})$$

$$= H(x_N) - H(x_N \mid x_1, \dots, x_{N-1})$$

$H = E(\log p) = \text{Entropy}$


$H_r = \text{Entropy Rate}$



Papers on Information Dynamics

- Dubnov, S. (2004). "Generalization of Spectral Flatness Measure for Non-Gaussian Linear Processes." *IEEE Signal Processing Letters*, 11(8), 698–701
- Dubnov, S. (2006). "Spectral Anticipations." *Computer Music Journal*, 30(2), 63–83.
- Abdallah, S., & Plumbley, M. (2010). Information dynamics: Patterns of expectation and surprise in the perception of music. *Connection Science*, 21(2-3), 89–117.
- IDyOM: <https://www.marcus-pearce.com/idyom/>
- Controlling Surprisal in Music Generation via Information Content Curve Matching: <https://arxiv.org/pdf/2408.06022>
- Estimating Musical Surprisal in Audio: <https://arxiv.org/pdf/2501.07474>
- Estimating Musical Surprisal from Audio In Autoregressive Diffusion Model Noise Spaces: : <https://arxiv.org/pdf/2508.05306>

Factor Oracle String Matching




Allauzen, Crochemore, Raffinot, 1999


Compression with COMPROR

Idea: Replace Entropy H with Coding Length C


$$IR(\text{letter}) = C(\text{letter}) - C(\text{letter} \mid \text{past sequence}) = | \text{Single letter encoder} | - | \text{Block encoder} | / \text{Block size}$$




Compror encoding scheme



- Compror code length is used as an approximation to Shannon / Kolmogorov Entropy
$$IR = \log(\text{size of alphabet}) - \# \text{ bits} / \text{block length}$$
- Search over all thresholds to choose a model that has highest IR




Multi-Agent Improvisation



Directed Information

- *Accounts for the intra-dynamics of parallel processes*
- *Related to causality*


$T_{y \rightarrow x}$



$$I_1(X^T \rightarrow Y^T) = \sum_{t=1}^T I(X^{t-1}; Y_t | Y_{<t}) = \sum_{t=1}^T \left[H(Y_t | Y_{<t}) - H(Y_t | Y_{<t}, X^{t-1}) \right].$$

Information Dynamics in Multi-Track Models

- Each **Voice is an agent** (human and machine)
- Improvisation: **Listening <-> Playing**
- **Information** passed between agents
- Creative interaction is measured by the amount of **information transfer** between man and the machine



Total Information Flow

Hypothesis: creative interaction maximizes the total information flow. Given two interactive signals X and Y

$$Flow = T_{X \rightarrow Y} + T_{Y \rightarrow X}$$

$$T_{X \rightarrow Y} + T_{Y \rightarrow X} = H(X|\bar{X}) + H(Y|\bar{Y}) - H(XY|\overline{XY})$$

Total information flow is given by the surprisal factor of X and Y and reducing the surprisal of the joint signal

Estimating Information Flow

- Calculating Entropy is hard
 - High dimensional audio data
 - Complex MIDI data
- Use pre-trained models (generative)


$$H(Y|X) = - \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x, y) \log \frac{p(x, y)}{p(x)}$$

$$H(Y|X) = E \left[-\log \frac{p(x, y)}{p(x)} \right] = E [-\log p(y|x)]$$

Algorithm 1 Calculate Information Flow

```
1: Input: A MIDI file with two tracks (XY)
2: Output: Total Information Flow
3:  $X, Y \leftarrow \text{SPLITTRACKS}(XY)$   $\triangleright$  Separate two tracks
4: for each  $T \in \{X, Y, XY\}$  do
5:   for  $i \leftarrow 0$  to  $\text{seq\_len}$  do
6:      $\text{start} \leftarrow T[i : \text{context}]$ 
7:      $\text{probs} \leftarrow \text{MTMT.TRANSFORMER}(\text{start})$ 
8:      $\text{cond\_entropy}[T][i] \leftarrow - \sum \text{probs} \times \log(\text{probs})$ 
9:   end for
10:   $\text{cond\_entropy}[T] \leftarrow \text{AVG}(\text{cond\_entropy}[T])$ 
11: end for
12:  $\text{InfoFlow} \leftarrow \text{cond\_entropy}[X] + \text{cond\_entropy}[Y] - \text{cond\_entropy}[XY]$ 
13: return  $\text{InfoFlow}$ 
```


Visualizing Musical Transformer



<https://storage.googleapis.com/nips-workshop-visualization/index.html>

Multitrack Music Transformer (MMT)

$$\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \quad \mathbf{x}_i = (x_i^{type}, x_i^{beat}, x_i^{position}, x_i^{pitch}, x_i^{duration}, x_i^{instrument})$$




(b)

(0, 0, 0, 0, 0, 0)
(1, 0, 0, 0, 0, 15)
(1, 0, 0, 0, 0, 36)
(1, 0, 0, 0, 0, 39)
(2, 0, 0, 0, 0, 0)
(3, 1, 1, 41, 15, 36)
(3, 1, 1, 65, 4, 39)
(3, 1, 1, 65, 17, 15)
(3, 1, 1, 68, 4, 39)
(3, 1, 1, 68, 17, 15)
(3, 1, 1, 73, 17, 15)
(3, 1, 13, 68, 4, 39)
(3, 1, 13, 73, 4, 39)
(3, 2, 1, 73, 12, 39)
(3, 2, 1, 77, 12, 39)
(4, 0, ..., 0, 0, 0)

Start of song
Instrument: accordion
Instrument: trombone
Instrument: brassy
Start of notes
Note: beat=1, position=1, pitch=E2, duration=48, instrument=trombone
Note: beat=1, position=1, pitch=E4, duration=12, instrument=brassy
Note: beat=1, position=1, pitch=E4, duration=72, instrument=accordion
Note: beat=1, position=1, pitch=G4, duration=12, instrument=brassy
Note: beat=1, position=1, pitch=G4, duration=72, instrument=accordion
Note: beat=1, position=1, pitch=C5, duration=72, instrument=accordion
Note: beat=1, position=13, pitch=G4, duration=12, instrument=brassy
Note: beat=1, position=13, pitch=C5, duration=12, instrument=brassy
Note: beat=2, position=1, pitch=C5, duration=36, instrument=brassy
Note: beat=2, position=1, pitch=E5, duration=36, instrument=brassy
...
End of song


(c)




Dong, Hao-Wen et al. "Multitrack Music Transformer." *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2022): 1-5.

Is Flow Happening in Real Music?

- Evaluated on MIDI data from Musescore
- Randomly juxtaposed tracks to form negative examples





Mozart –
String Quartet K. 387

Violin, Staff

Violin, Staff-1


Viola, Staff-2

Violoncello, Staff-3

♩ = 120

3

Explaining listening to Generative Music Audio



DEFINITIONS	h = mean NLL, Eq. 1
PMus-Duo $h(\hat{O} \bar{M}) - h(\hat{O} M)$	Context gain · M1
PMus-Ctx $h(\hat{O} \emptyset) - h(\hat{O} M)$	Duo grounding · M4
PMus-Self $h(\hat{O} B) - h(\hat{O} M)$	Self-continuation gain · M6
PMus-Partner+ $h(\hat{O} \text{mix}(A, \bar{B})) - h(\hat{O} M)$	Decoy-controlled partner gain · M7
PMus-Partner $h(\hat{O} A) - h(\hat{O} M)$	Partner gain · M5

NOTATION	
$\hat{O}$	generated clip under evaluation
A	true own past (player A)
B	true partner's past (player B)
$M = \text{mix}(A, B)$	true past mix
$\bar{B}$	decoy partner's past, own past kept (mean of 8)
$\bar{M}$	decoy past mix, fully foreign duo (mean of 8)
$\emptyset$	no context

Predictive musical information and its agreement with listeners.

Left: share of judged pairs where each measure agrees with human choice,


Right: Five Measures with NLL contrasts

MusicGen
InspireMusic
YuE
Magenta RT
ACE-Step

Distributed intelligence



Multi-agent interaction: task reward and information



Standard MARL

Agents learn policies for shared or competing task rewards.

Empowerment probes

Hypothetical actions quantify control over one's own future state.

Nash / iterative water-filling

Best-response updates of probing distributions reach a fixed point.

What empowerment adds to a Vicsek flock



Vicsek baseline

Local alignment leads toward one shared heading. Average empowerment declines.

With empowerment control


Agents also adjust their turning to preserve influence over their own future headings.

Higher average empowerment; two opposed, organized streams instead of global alignment.

[Play the flock simulation](#)



Toy DI pilot: four coupled activations



Independent noise

Untrained baseline. Ring score 0.086.

Intrinsic only

CEM selects activity and diversity.
Ring score 0.079.

Intrinsic + DI proxy

CEM adds a ring bonus. Ring score 0.443.

Next: DI-Walker tests whether information relates to control under limb faults.

Communication for fault robustness

No peer sensing versus delayed peer-force messages during weak central-control dropout and limb-4 slip

Permanent limb loss

$$e_{4,t} = 0$$

Realized force is removed after the event.

Intermittent slip

$$e_{4,t} = 0.2$$

Force authority falls during selected windows.

Sensor loss

A selected message channel is removed or zeroed.

Weak-CC dropout

The shared correction signal is unavailable.

Compound failure

Limb loss, sensing loss, and weak-CC dropout co-occur.

<https://arxiv.org/abs/2609.27580>

<https://github.com/sdubnov/DI-Walker>

V4.3d Robust Walker: task=chirp seed=9105

Large-layout replay: intermittent central control + intermittent L4 slip

step: 000

CC: on

L4 slip: recovered

Capacity / No Peer Sensing

Sensor-Comm



Late path error: 1.414

— Target
— Body path



Late path error: 0.363

— Target
— Body path

Animated in slide show

Paper connection - DI-Walker v4.3 failure replay

Distributed action pathway

Realized force is available to the next controller update




Own-Sensor: $q_{i,t} = q_{i,t}^{local} + w_i s_{i,t}$

Peer-Sensor: $q_{i,t} = q_{i,t}^{local} + \sum_{j \neq i} w_{ij} s_{j,t}$


$G_{act} = \mathbb{E}[\log p(U_t|C_t, M_t) - \log p(U_t|C_t)] / \ln 2$




Action-directed information: G_{act}



Compound failure: L4 loss followed by weak-CC dropout



Functional robustness under disturbance




Action-directed Information:



- history = 4, predictor receives the window of 5 samples
- Full-precision sensor messages, without quantization.
- A context-only MLP versus a context-plus-message MLP.

Predictor-only results

Peer messages retain large action-predictive value under compound failure, even at coarse resolution



Predictor-only message quantization: IT-PAC-style action information



$$G_{\text{act}} \approx 3.3\text{--}3.6 \text{ bits/action}$$

stable across finite message levels

$$\Delta_{\text{shuffle}} : 0.02 \rightarrow 0.59 \text{ bits/action}$$

alignment detail improves with resolution

$$\eta_{\text{act}} = G_{\text{act}} / R_{\text{nominal}}$$

ratio declines as nominal rate grows

Error bars show one standard deviation across three predictor seeds; efficiency is descriptive, not channel capacity.

Main Concepts:

Information flowing over time is relevant when the transmitted signal depends on past channel outputs or feedback. This is especially important for **interactive, multi-agent, and control-oriented communication**. The toy experiment shows an increase in **action-directed information** for distributed messaging and an improvement in fault robustness even at lower bit rates.

Empowerment measures how much control an agent has over its future observations or states:

$$\mathcal{E}(s_t) = \max_{p(a_{t:t+H})} I(A_{t:t+H}; S_{t+H} \mid S_t = s_t).$$

It is essentially an action-to-future-state channel capacity: how many distinguishable future outcomes the agent can create.

Directed information (DI) measures causal information flow through time:

$$I(X^n \rightarrow Y^n) = \sum_{t=1}^n I(X^t; Y_t \mid Y^{t-1}).$$

For closed-loop systems, empowerment can be defined using directed information from actions to future observations:

$$\mathcal{E}_{\text{closed}} = \max_{\pi} I(A^n \rightarrow O^n),$$

where π denotes an admissible feedback policy.

(Abel et al. Plasticity as the Mirror of Empowerment <https://arxiv.org/html/2505.10361v2>)

Relevance to Multi-Agent Interaction and Robustness:

In a distributed agentic network, there are at least three relevant information flows:

local observation → semantic message → other agent's action → future shared state.

They can be measured differently:

- **Observation-to-message DI:** how much causally relevant information is extracted and transmitted.
- **Message-to-action DI:** how much the message changes the receiver's decisions.
- **Action-to-future-state empowerment:** how much control the resulting actions provide over the environment.
- **Team empowerment:** how much influence the distributed group has over future collective states.

Directed information is useful for measuring whether communication has a causal effect:

$$I(M^n \rightarrow A^n) = \sum_{t=1}^n I(M^t; A_t | A^{t-1}),$$

where M^n is the message sequence and A^n is the action sequence.

A robust message should have a high **causal influence** on the correct action that remains stable when messages are delayed, corrupted, dropped, or partially missing.



mdpi.com/si/271865



**Open for
Submission!**

- **Special Issue:** Information Dynamics and Control in Generative Cognitive Systems
- **Guest Editor:** Prof. Shlomo Dubnov, Dr. Stas Tiomkin, Dr. Stefan Lattner
- **Submission deadline:** 30 April 2027

Special Issue Information:

This Special Issue focuses on such measures—like predictive information, directed information, and information rate—and their role in systems that learn, act, and generate structured behaviors (e.g., in language, music, or autonomous agents). We invite original contributions, as well as comprehensive surveys, including, but not limited to, the following areas:


- Predictive information and anticipatory structure in dynamical systems;
- Information content (IC), surprisal, and time-local information measures ;
- Information rate (IR) and temporal constraints on perception, action, and generation;
- Directed information, causal information flow, and influence between processes;
- Information-to-go, controllable information production, and information-based control ;
- Information dynamics in generative models (e.g., music, language, symbolic sequences);
- Surprisal shaping, information curve matching, and controllable generation;
- Action–perception loops analyzed through dynamic information measures;
- Intrinsic motivation and exploration driven by predictive or directed information;
- New information measures explicitly designed for temporal, generative, or interactive systems.

Extra Technical Slides


Information in Control and Perception-Action

Actions change the world, which changes what the agent senses next.

IT-PAC describes this closed loop in a common information-theoretic language.



The next state becomes the next observation




$$V^\pi(s_t, a_t) \quad \mathcal{I}^\pi(s_t, a_t)$$

Future reward or
control performance

Information processed
over the future state-
action sequence


$$\Delta I_t = \log \frac{p(s_{t+1} | s_t, a_t)}{\hat{p}(s_{t+1})} + \log \frac{\pi(a_{t+1} | s_{t+1})}{\hat{\pi}(a_{t+1})}$$

Fig. 2 Policies resulting from the trade-off between value-to-go and information-to-go



Implemented perception-action causal ordering

The message is computed from delayed realized peer force before the receiving action is chosen



$$F_{-i,t-1} \longrightarrow M_{i,t} \longrightarrow U_{i,t} \longrightarrow Z_{t+1}$$

$$M_{i,t} = Q_B(F_{-i,t-1})$$

IT-PAC estimator: message value at the action interface

Two predictors share the same ordinary context, while only one receives the delayed sensor-message history

$$C_{i,t}^{(k)} = (C_{i,t-k}, \dots, C_{i,t}), \quad M_{i,t}^{(k)} = (M_{i,t-k}, \dots, M_{i,t})$$

Context-only baseline

$$q_0 \left(u_{i,t} \mid C_{i,t}^{(k)} \right)$$

same target actions



Context plus true message

$$q_1 \left(u_{i,t} \mid C_{i,t}^{(k)}, M_{i,t}^{(k)} \right)$$

Action-predictive gain

$$G_{\text{act}} = \frac{\mathbb{E}[\text{NLL}(q_0)] - \mathbb{E}[\text{NLL}(q_1)]}{\ln 2}$$

True-message advantage

$$\Delta_{\text{shuffle}} = \frac{\mathbb{E}[\text{NLL}(q_{1,\text{shuf}})] - \mathbb{E}[\text{NLL}(q_1)]}{\ln 2}$$

$C_{i,t}$: ordinary context, $M_{i,t}$: sensor message, k : history depth, q_0, q_1 : conditional action densities, $\text{NLL} = -\log q(u)$

* We trained an MLP predictor for next action with and without past messages