Species Field Theory: Discovery of Latent Ecological Structure from Community Time Series

XINGJI CUI[1,*]

[1]Xi'an Jiaotong University, China
[*]Correspondence: cuixingji@gmail.com

Abstract. Ecological abundance time series are shaped not only by interactions among species, but also by broader community-level dynamics such as hidden resources and shared ecological constraints. We introduce Species Field Theory (SFT), a field-based framework for discovering latent ecological structure from abundance time series. SFT recovered directed interactions in a six-species Lotka–Volterra system, achieving a mean signed Spearman correlation of 0.887 across five random seeds. In a Huisman resource-competition system, SFT latent states aligned with hidden resource dynamics. These results suggest that interaction recovery is a special case of broader latent ecological structure discovery.

Key words: Species Field Theory; ecological time series; latent ecological structure; ecological dynamics; field-based representation; machine learning.

Introduction

Ecological abundance time series reflect underlying ecosystem structures. These structures may take many forms—direct species–species interactions (Lotka, 1925; Volterra, 1926; May, 1972), limiting resources (Tilman, 1982; Huisman and Weissing, 1999), environmental forcing or seasonality (Moran, 1953; Vasseur and Yodzis, 2004), or latent community-level constraints—and a learning framework should not assume in advance that they reduce to a pairwise interaction matrix.

Existing methods address parts of this problem. State-space models capture unobserved processes and observation noise, empirical dynamic modeling reconstructs nonlinear system states for forecasting and causality, and neural dynamical models such as Neural ODEs and Latent ODEs learn continuous-time hidden dynamics from data (Holmes et al., 2012; Sugihara et al., 2012; Chen et al., 2018; Rubanova et al., 2019). However, these approaches do not by themselves specify how learned latent states should be interpreted: as pairwise interactions, resource-mediated couplings, external drivers, or broader community-level structure. Structure identification, rather than prediction alone, is our primary objective.

We introduce Species Field Theory (SFT), a field-based framework for latent ecological structure discovery from abundance time series. Each species is represented by a dynamic field vector with both static and community-dependent components. We demonstrate SFT on two contrasting simulated systems: a six-species Lotka–Volterra model with an explicit directed interaction matrix, and a Huisman resource-competition model where species are coupled indirectly through limiting resources. Together, these contrasting cases illustrate our central claim: pairwise interaction recovery is one special case of latent ecological structure discovery.

Methods

Problem formulation

Let $\mathbf{x}_t \in \mathrm{R}_+^N$ denote the abundance vector of N observed species at time t, and let $\mathbf{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_T\}$ be the observed community time series. We model community change through the log-abundance increment

$$\Delta \mathbf{y}_t = \log(\mathbf{x}_{t+1} + \varepsilon) - \log(\mathbf{x}_t + \varepsilon), \qquad t = 1, \ldots, T-1,$$

where $\varepsilon > 0$ is a small constant used for numerical stability.

Our goal is not only to predict $\Delta \mathbf{y}_t$, but to learn a latent ecological field from which structure can be read out and compared. This structure is not assumed to be pairwise in advance: depending on the system, it may correspond to direct species interactions, resource-mediated coupling, external forcing, or broader community-level constraints.

Species field representation

SFT represents each species by a time-dependent field vector rather than a fixed embedding. A community encoder maps

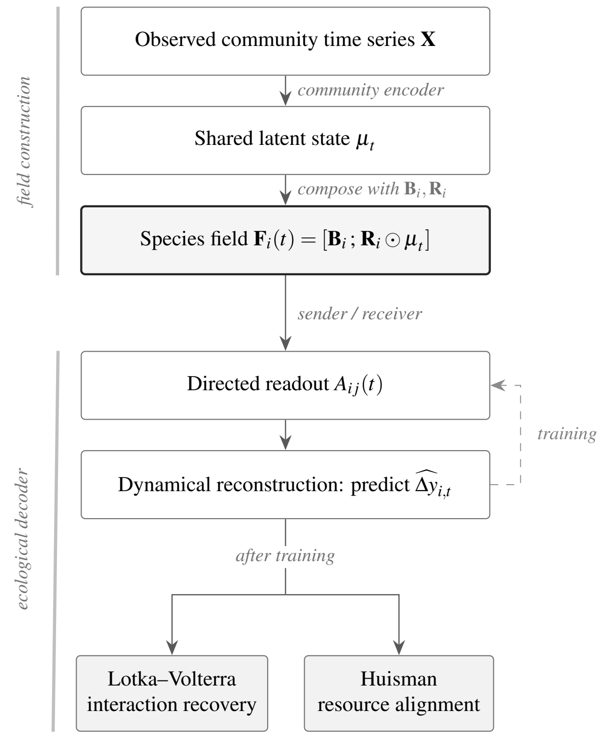

Figure 1. Overview of the SFT framework. The community time series **X** is encoded into a shared latent ecological state μ_t, which is composed with a static species identity $\mathbf{B}_i$ and a species-specific response $\mathbf{R}_i$ to form the species field $\mathbf{F}_i(t)$ (centerpiece). Sender/receiver projections of $\mathbf{F}_i(t)$ produce a directed interaction-like readout $A_{ij}(t)$, which combines with observed abundances to predict species-level increments $\widehat{\Delta y}_{i,t}$ as a training signal. After training, the same field representation supports two ecological readouts: pairwise interaction recovery for the Lotka–Volterra system, and latent resource alignment for the Huisman resource-competition system.

the observed abundance sequence to a time-indexed latent ecological state $\mu_t \in \mathbb{R}^{d_{\text{latent}}}$, shared across all species. Each species combines a static identity component with a species-specific response to this shared state:

$$\mathbf{F}_i(t) = [\mathbf{B}_i \,;\, \mathbf{R}_i \odot \mu_t],$$

where $[\,;\,]$ denotes vector concatenation and $\odot$ denotes the element-wise product. Here $\mathbf{B}_i \in \mathbb{R}^{d_{\text{static}}}$ is a static species identity component and $\mathbf{R}_i \in \mathbb{R}^{d_{\text{latent}}}$ is a species-specific response vector, so $\mathbf{F}_i(t) \in \mathbb{R}^{d_{\text{static}}+d_{\text{latent}}}$. The field $\mathbf{F}_i(t)$ thus separates persistent species identity from the way species i expresses the community-level ecological state at time t.

Directed structure readout

To examine directional ecological structure, SFT projects each species field into separate sender and receiver representations:

$$\mathbf{s}_i(t) = \mathbf{W}_{\text{send}}\mathbf{F}_i(t), \qquad \mathbf{r}_i(t) = \mathbf{W}_{\text{recv}}\mathbf{F}_i(t),$$

where $\mathbf{W}_{\text{send}}, \mathbf{W}_{\text{recv}} \in \mathbb{R}^{d_{\text{readout}} \times (d_{\text{static}}+d_{\text{latent}})}$ and $\mathbf{s}_i(t), \mathbf{r}_i(t) \in \mathbb{R}^{d_{\text{readout}}}$. A directed interaction-like score is then defined by

$$A_{ij}(t) = \frac{\langle \mathbf{r}_i(t), \mathbf{s}_j(t) \rangle}{\sqrt{d_{\text{readout}}}}, \qquad i \neq j,$$

analogous to the scaled dot-product form used in attention models. By convention, $A_{ij}(t)$ denotes the learned influence score of species j on species i; self-effects are excluded from this readout and modeled separately. Because the sender and receiver projections are distinct, $A_{ij}(t) \neq A_{ji}(t)$ in general.

The interpretation of $A_{ij}(t)$ depends on the underlying system. In systems with explicit pairwise interactions, it can be compared directly with the true interaction matrix. In systems governed by indirect coupling, such as resource-mediated competition, $A_{ij}(t)$ should instead be read as an interaction-like diagnostic of the learned field rather than as a ground-truth pairwise mechanism.

Dynamical reconstruction

The field representation is learned by requiring it to reconstruct observed community change. Given the directed readout $A_{ij}(t)$, SFT predicts the log-abundance increment of species i as a sum of a species-specific self-dynamics term and interaction-like contributions from the rest of the community:

$$\widehat{\Delta y}_{i,t} = g_i(x_i(t)) + \sum_{j \neq i} A_{ij}(t)\, x_j(t),$$

where $g_i(x_i(t)) = r_i + d_i x_i(t)$ is a linear self-dynamics term with species-specific parameters $r_i, d_i \in \mathbb{R}$. The second term aggregates interaction-like contributions read from the learned field. The model is trained by minimizing the discrepancy between $\widehat{\Delta y}_{i,t}$ and the observed increment $\Delta y_{i,t}$, with mild regularization on μ_t to encourage stable and smooth latent trajectories.

Prediction here serves as a learning signal, not as the primary scientific target. The target is the latent field $\mathbf{F}_i(t)$ and the ecological structures that can be read from it after training.

Simulated ecological systems

We evaluated SFT on two simulated systems chosen to represent contrasting forms of ecological structure.

The first was a six-species Lotka–Volterra system in which community dynamics are generated by an explicit directed pairwise interaction matrix. This setting provides a ground truth against which the learned readout $A_{ij}(t)$ can be compared directly.

The second was a Huisman resource-competition system with six species and five limiting resources (Huisman and Weissing, 1999). Species are not coupled through an explicit pairwise interaction matrix; instead, they interact indirectly by consuming shared resources. The relevant ecological structure here is resource-mediated rather than directly pairwise, and no LV-style direct pairwise ground truth exists.

Together, these two systems allow us to examine whether the same field-based representation supports different semantic readouts: directed interaction recovery when the true structure is pairwise, and latent resource alignment when the true structure is mediated by hidden resources.

Structure readout and evaluation

After training, we evaluated the learned field using a system-specific criterion for each simulated system.

For the Lotka–Volterra system, where the true structure is an explicit directed interaction matrix, we summarized the learned directed readout $A_{ij}(t)$ by averaging over the validation period and compared its off-diagonal entries with the true interaction matrix. The primary metric was the signed Spearman correlation between learned and true off-diagonal entries; we also reported sign accuracy to assess whether positive and negative effects were recovered correctly.

For the Huisman system, direct pairwise recovery was not treated as the primary target, because the system is generated by resource-mediated coupling rather than by an explicit interaction matrix. We instead asked whether the learned latent ecological state μ_t was aligned with the hidden resource dynamics. For each simulated resource trajectory, we computed the absolute Pearson correlation with every latent dimension of μ_t over the validation period and recorded the maximum across dimensions; we refer to this as the *latent resource alignment* of that resource. This score is used as a diagnostic of resource-related information in the latent field, rather than as a claim of full resource recovery.

The evaluation criterion thus depends on the ecological mechanism: directed interaction recovery for the pairwise Lotka–Volterra system, and latent resource alignment for the resource-mediated Huisman system. Prediction error was monitored during training but was not the primary scientific evaluation; the primary evaluation was whether the learned field supported the appropriate ecological interpretation in each setting.

Results

SFT recovers directed interactions in Lotka–Volterra dynamics

The main structure readouts are shown in Fig. 2. In the six-species Lotka–Volterra system, the generative ecological structure is an explicit directed interaction matrix. This provides a setting in which the directed readout learned by SFT can be evaluated directly against the known mechanism.

Across five random seeds, SFT recovered the directed pairwise structure from abundance time series alone. The learned interaction-like readout achieved a mean signed Spearman correlation of 0.887 with the true off-diagonal interaction matrix, with a standard deviation of 0.013 across seeds. Sign accuracy was 0.90 in all five runs, indicating that the model consistently recovered most positive and negative off-diagonal interaction signs.

These results suggest that when the underlying ecological mechanism is genuinely pairwise and directed, the learned field can support recovery of the corresponding interaction matrix. Pairwise interaction recovery is therefore one interpretable readout of the SFT field when such a readout is mechanistically appropriate.

SFT latent states align with hidden resources in resource competition

The Huisman resource-competition system presents a different ecological setting. Species are coupled indirectly through shared limiting resources rather than through an explicit Lotka–Volterra-style interaction matrix. Direct pairwise interaction recovery is therefore not the natural target for this system.

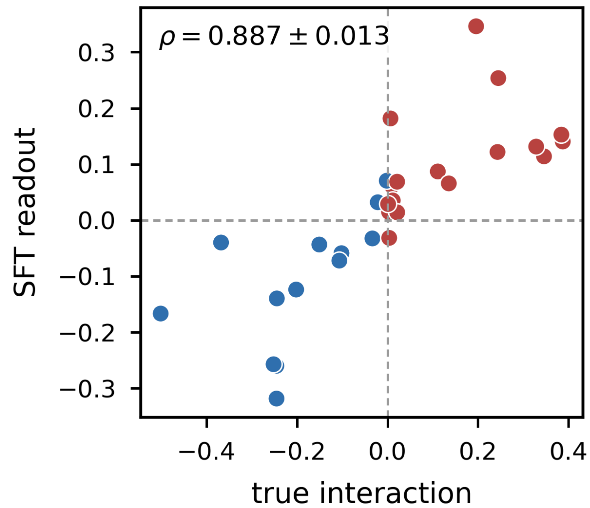

(a)

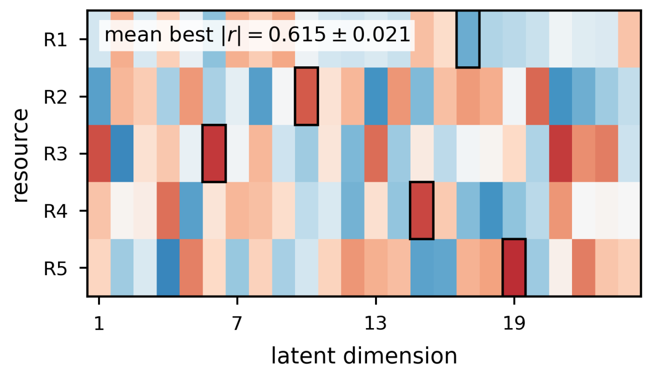

(b)

Figure 2. Main quantitative readouts of SFT. (a) Off-diagonal Lotka–Volterra interaction coefficients plotted against the corresponding SFT readout values for a representative run. Across five random seeds, the mean signed Spearman correlation was $\rho = 0.887 \pm 0.013$ and sign accuracy was 0.90 in all runs. (b) Correlations between hidden Huisman resource trajectories and latent dimensions of μ_t for a representative run. Black boxes mark the best-aligned latent dimension for each resource; across three random seeds, the mean best absolute correlation was 0.615 ± 0.021. Resource alignment is interpreted as a diagnostic of resource-related information in the latent field, not as full resource recovery.

In this setting, the directed readout did not provide a strong direct interaction interpretation. However, the learned latent ecological state showed reproducible alignment with the hidden resource dynamics. Across three random seeds, the mean latent resource alignment was 0.615, with a standard deviation of 0.021. This indicates that the latent field captured resource-related variation even though the model was trained only from species abundance trajectories.

Together, the two systems illustrate a distinction between pairwise interaction recovery and broader latent ecological structure discovery. In the Lotka–Volterra system, the appropriate semantic readout is a directed interaction matrix. In the Huisman system, the relevant structure is resource-mediated and is more naturally reflected in the latent ecological state.

These results support the central claim that pairwise interaction recovery is one special case of latent ecological structure discovery.

Discussion

This study presents SFT as a minimal field-based framework for latent ecological structure discovery from abundance time series. Its main contribution is not a new predictor of ecological trajectories, but a shift in representation: ecological structure is treated not as a fixed pairwise interaction matrix, but as something that may be expressed through a learned ecological field. The interaction matrix is therefore not the primary object of the theory; it is one possible readout of the field when the underlying mechanism is genuinely pairwise.

This field-first view is motivated by several theoretical perspectives. From dynamical systems theory, ecological abundance time series are observations of an underlying system whose state may include population densities, resources, environmental conditions, and other hidden variables. Takens' embedding theorem suggests that, under suitable conditions, delay-coordinate observations can reconstruct the geometry of an underlying attractor up to a smooth transformation (Takens, 1981; Sauer et al., 1991). We therefore view the latent ecological state μ_t in SFT as a learned coordinate system for aspects of the reconstructed community state, rather than as a directly observed environmental variable.

A related representational intuition comes from Koopman operator theory. Nonlinear dynamical systems can be studied through the evolution of observables in a lifted function space, where nonlinear dynamics may reveal more structured representations (Koopman, 1931; Mezic, 2005). SFT is not a Koopman model, but it follows a similar logic: rather than analyzing species abundances only in their original coordinate system, it learns field-valued observables $\mathbf{F}_i(t)$ from which ecological structure can be read. The sender–receiver readout is therefore best understood as a finite-dimensional structural readout from the learned field. In the Lotka–Volterra case, this readout can be interpreted as a directed interaction matrix; in resource-mediated systems, the same field may instead organize information about hidden resource dynamics.

The field construction also has an ecological interpretation through niche theory. Species are not only nodes in an interaction network; they are differentiated by how they respond to shared environmental, resource, and community-level axes. Classical niche theory views species coexistence and competition through positions and overlaps in resource or environmental space (Hutchinson, 1957; MacArthur and Levins, 1967; Tilman, 1982). In SFT, the shared latent state μ_t represents community-level ecological context, while the species-specific response vector $\mathbf{R}_i$ determines how species i expresses that context. Thus, $\mathbf{R}_i \odot \mu_t$ can be interpreted as a learned niche-response expression: the same ecological state is shared by the community, but different species respond to it differently.

The two experiments illustrate why this distinction matters. In the Lotka–Volterra system, the true mechanism is explicitly pairwise and directed. The learned field can therefore be projected into a directed interaction readout, which recovers the true interaction structure with high signed rank correlation. This result should not be interpreted as evidence that all ecological systems are pairwise; rather, it shows that when pairwise interaction is the correct ecological semantics, the SFT field contains enough directional structure for that matrix to emerge as a valid readout.

The Huisman resource-competition system gives the complementary lesson. Species are coupled through hidden limiting resources rather than through a direct Lotka–Volterra-style interaction matrix. A weak direct interaction readout is therefore not merely a negative result; it indicates that a mismatched semantic projection is being applied to a resource-mediated system. The reproducible alignment between the learned latent state and hidden resource dynamics suggests that SFT can capture ecological structure that is not naturally expressed as pairwise species–species coefficients.

These results motivate a more unified field-first view of ecological structure discovery. The goal should not be to assign a separate decoder to every dataset or ecological mechanism. Instead, future work should introduce an ecological operator layer between field construction and semantic interpretation. Such an operator would act on the learned ecological field and determine how structure is expressed. Pairwise interactions, resource-mediated constraints, external forcing, seasonal variation, and higher-order community constraints could then be treated as different structural expressions of the same underlying field, rather than as unrelated modeling targets.

Future development of SFT should move in three directions. First, the present species-indexed and time-indexed discrete field should be generalized toward a continuous ecological field over trait, niche, phylogenetic, or learned latent coordinates. Second, a stricter mathematical foundation is needed to connect Takens-style state reconstruction, Koopman-style observable dynamics, and ecological field construction. Third, ecological operators should be developed to provide a unified bridge between the learned field and its ecological interpretations. Applying this framework to real communities will also require explicit treatment of external drivers, seasonality, migration, observation noise, and hidden community structure.

Overall, these results support the central claim that pairwise interaction recovery is one special case of latent ecological structure discovery. SFT provides an initial field-based representation in which abundance time series can be lifted into an ecological state space, species can express shared community states through species-specific responses, and different ecological mechanisms can become readable as structures of a common ecological field.

Acknowledgements

The author thanks Suzhou High School SIP and Mr. Jun Zhao for support. The author is grateful to Haoyuan Sun, Yanhong Pan, Yifan Situ, Tianyou Zhang, and his family for encouragement.

Code and data availability

Code and processed data are available from the author upon reasonable request and will be made publicly available in a future version.

References

Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, volume 31.

Holmes, E. E., Ward, E. J., and Wills, K. (2012). Marss: Multivariate autoregressive state-space models for analyzing time-series data. *The R Journal*, 4(1):11–19.

Huisman, J. and Weissing, F. J. (1999). Biodiversity of plankton by species oscillations and chaos. *Nature*, 402(6760):407–410.

Hutchinson, G. E. (1957). Concluding remarks. In *Cold Spring Harbor Symposia on Quantitative Biology*, volume 22, pages 415–427.

Koopman, B. O. (1931). Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, 17(5):315–318.

Lotka, A. J. (1925). *Elements of Physical Biology*. Williams and Wilkins, Baltimore, Maryland, USA.

MacArthur, R. and Levins, R. (1967). The limiting similarity, convergence, and divergence of coexisting species. *The American Naturalist*, 101(921):377–385.

May, R. M. (1972). Will a large complex system be stable? *Nature*, 238(5364):413–414.

Mezic, I. (2005). Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41:309–325.

Moran, P. A. P. (1953). The statistical analysis of the canadian lynx cycle. ii. synchronization and meteorology. *Australian Journal of Zoology*, 1(3):291–298.

Rubanova, Y., Chen, R. T. Q., and Duvenaud, D. K. (2019). Latent ordinary differential equations for irregularly-sampled time series. In *Advances in Neural Information Processing Systems*, volume 32.

Sauer, T., Yorke, J. A., and Casdagli, M. (1991). Embedology. *Journal of Statistical Physics*, 65:579–616.

Sugihara, G., May, R., Ye, H., Hsieh, C.-h., Deyle, E., Fogarty, M., and Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, 338(6106):496–500.

Takens, F. (1981). Detecting strange attractors in turbulence. In Rand, D. and Young, L.-S., editors, *Dynamical Systems and Turbulence, Warwick 1980*, volume 898 of *Lecture Notes in Mathematics*, pages 366–381. Springer, Berlin.

Tilman, D. (1982). *Resource Competition and Community Structure*. Princeton University Press, Princeton, New Jersey, USA.

Vasseur, D. A. and Yodzis, P. (2004). The color of environmental noise. *Ecology*, 85(4):1146–1152.

Volterra, V. (1926). Fluctuations in the abundance of a species considered mathematically. *Nature*, 118:558–560.