

A Unified and Predictive Measure of Functional Diversity

Adji Bousso Dieng^{1, 3, *} and Amey P. Pasarkar^{2, 3}

¹Department of Computer Science, Princeton University

²Lewis-Sigler Institute For Integrative Genomics, Princeton University

³Vertaix

*Corresponding author: adji@princeton.edu

September 19, 2025

Abstract

Despite the critical role of functional diversity (FD) in understanding ecological systems and processes, its robust quantification remains a significant challenge. A long-held view in the field is that it is not possible to capture its three facets—functional richness, functional divergence, and functional evenness—in a single index. This perspective has prompted recent proposals for FD measurement to use three separate indices, one for each aspect. Here, we challenge this paradigm by demonstrating that the probability-weighted Vendi Score (pVS), first introduced by [Friedman and Dieng \(2023\)](#), can serve as a powerful functional diversity index that can capture its three facets. We adapt pVS to functional ecology by defining it as the exponential of the Rényi entropy of the eigenvalues of the abundance-weighted trait similarity matrix. This formulation allows pVS to be applicable at any biological level. It can be defined at the species level, at which most existing FD metrics are defined, and at the individual level to naturally incorporate intraspecific trait variation (ITV) when detailed data are available. We theoretically and empirically demonstrate the robustness of pVS. We first mathematically prove it satisfies several essential desiderata for FD metrics, including invariance to functional redundancy, set monotonicity, distance monotonicity, and concavity. We then show that pVS consistently exhibits the expected ground-truth behavior on simulated ecosystem scenarios under which many FD metrics fail. By integrating abundances and trait similarities within a single, theoretically sound framework, pVS provides a generally applicable index for ecology.

Keywords: Diversity, Functional Diversity, Intraspecific Trait Variation, Biodiversity, Ecosystem Function, Resilience, Theoretical Ecology, Vendi Scoring

1 Introduction

The robust and comprehensive quantification of functional diversity (FD) has remained a persistent challenge in ecology. FD is commonly conceptualized through three distinct facets: functional richness, functional divergence, and functional evenness ([Villéger et al., 2008](#); [Schmera et al., 2023](#)). Functional richness is the

number of distinct functional types or the volume of functional trait space occupied by the community. It reflects the variety of functional roles present. Functional divergence on the other hand describes how functionally dissimilar individuals or species are from each other and how evenly they are spread across the functional trait space. Finally, functional evenness reflects the uniformity of the distribution of abundances across the different functional types or positions in trait space. High functional evenness implies that functional roles are equally represented in terms of abundance.

A long-held view in the field posits that it is inherently difficult, if not impossible, to capture these three distinct facets within a single, unified index (Villéger et al., 2008; Mouchet et al., 2010). This difficulty arises because each facet emphasizes different aspects of trait distribution. This perspective has historically led to the development of numerous, often disparate, metrics (Villéger et al., 2008; Mouchet et al., 2010; Schmera et al., 2023), with even very recent proposals continuing this trend by introducing separate indices tailored to each aspect (e.g., NJ Trees Richness, NJ Trees Divergence, and NJ Trees Evenness (Cardoso et al., 2024)). Furthermore, many existing metrics often overlook or struggle to adequately incorporate intraspecific trait variation (ITV), which is increasingly recognized as essential for a more nuanced and accurate quantification of FD (Palacio et al., 2024). These challenges underscore the urgent need for a new generation of FD metrics that can provide a more accurate, holistic, and theoretically sound assessment of ecological systems.

This paper proposes a solution to the challenges raised above. We demonstrate that the probability-weighted Vendi Score (pVS), first introduced by Friedman and Dieng (2023), can serve as a flexible and accurate FD index. We adapt pVS to functional ecology by defining it as the exponential of the Rényi entropy of the eigenvalues of the abundance-weighted trait similarity matrix. This formulation allows pVS to be applicable regardless of the unit of analysis. In particular, pVS can be defined at the individual level, allowing it to naturally incorporate intraspecific trait variation (ITV). This direct accounting for individual-level differences provides a more ecologically realistic and fine-grained assessment of FD than metrics relying solely on species-average traits, which often overlook crucial within-species variation. pVS can also be defined at the species level when individual-level data aren't available.

We mathematically prove that pVS rigorously satisfies several essential theoretical criteria for robust FD metrics. Notably, its adherence to properties such as set monotonicity (reflecting richness), distance monotonicity (reflecting divergence), and continuity in abundance (reflecting evenness) underscores its ability to holistically capture the multifaceted nature of FD. Furthermore, we demonstrate that pVS consistently exhibits the expected behavior across simulated ecosystem scenarios where ground-truth is known and where many existing FD metrics fail.

2 Theory

Consider an ecological community (C), which is composed of a collection of N observed biological individuals ($\mathbf{x}_1, \dots, \mathbf{x}_N$). Each of these individuals is characterized by a unique set of functional traits, and we denote by \mathbf{t}_i the traits of \mathbf{x}_i . These traits, such as body mass, leaf area, or metabolic rate, are measurable characteristics

that define an organism's ecological role and influence its performance within the ecosystem. In our framework, each individual \mathbf{x}_i is also associated with an abundance a_i , which quantifies its presence or contribution within the community. This abundance can take various forms: it might be a biomass measure (e.g., the total mass of a fish), a cover measure (e.g., the percentage of ground covered by a plant), or even a measure reflecting energy use or metabolic rate. Denote by p_i the relative abundance of \mathbf{x}_i , calculated as the proportion of an individual's abundance relative to the total abundance of all individuals in the community, ensuring that all relative abundances sum to one.

Once the traits and relative abundances of the members of C are defined, we construct a similarity matrix $\mathbf{K} \in \mathbb{R}^{N \times N}$ based on the pairwise trait distances between members. This matrix is generated using a positive semi-definite kernel similarity function $k(\cdot, \cdot)$ that is such that $k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_j, \mathbf{x}_i)$ and $k(\mathbf{x}_i, \mathbf{x}_i) = 1$. These properties of the kernel are reflected in the similarity matrix \mathbf{K} which is defined such that $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ and $K_{ii} = 1$. Note, the similarity function can be chosen depending on the application. For example, for quantitative traits, the Radial Basis Function (RBF) kernel with Euclidean distance is a common choice. Cosine similarity is another popular choice. For communities with a mix of quantitative and categorical traits, the Gower distance, or its generalizations (Pavoine et al., 2009; D'Orazio, 2021), can be used to build the similarity matrix. Denote by \mathbf{K}_p the relative abundance-weighted similarity matrix,

$$\mathbf{K}_p = \text{diag}(\sqrt{\mathbf{p}}) \mathbf{K} \text{diag}(\sqrt{\mathbf{p}}) \quad (1)$$

where $\text{diag}(\sqrt{\mathbf{p}})$ is the diagonal matrix whose diagonal entries correspond to the square root of the relative abundances $\sqrt{\mathbf{p}} = (\sqrt{p_1}, \dots, \sqrt{p_N})$. This matrix \mathbf{K}_p effectively combines trait similarities with relative abundances. Let $\lambda_1, \dots, \lambda_N$ denote the eigenvalues of \mathbf{K}_p . They reflect both functional similarity and abundance.

Applying the probability-weighted Vendi Score (pVS) (Friedman and Dieng, 2023) to functional ecology, we define the functional diversity of C as the exponential of the Renyi entropy of order q of the eigenvalues,

$$\text{pVS}(C, k, \mathbf{p}, q) = \exp\left(\frac{1}{1-q} \log \sum_{i=1}^N \lambda_i^q\right) \quad (2)$$

where we use the convention $0^0 = 0$ and $0 \log 0 = 0$. This definition of pVS can also be applied when the unit of analysis is a species, by simply defining the abundances and the traits at the species level.

The order q in Eq. 2 determines the sensitivity of pVS to trait variations and abundance distribution within the community. For $q = 0$, pVS quantifies the effective number of functionally distinct groups within the community, essentially simply counting how many truly unique functional roles are present. When $0 < q < 1$, pVS is sensitive to rare functional types or individuals with unique trait combinations, even if they have low abundances, thereby giving more weight to the less dominant functions in the community. Conversely, when $q > 1$, pVS increasingly emphasizes dominant functional types or individuals with high abundance and common trait combinations, with pVS accounting for only the most prevalent functional role

in the community as $q \rightarrow \infty$. Finally, when $q = 1$ (the limit as $q \rightarrow 1$), pVS corresponds to the Shannon entropy of the eigenvalues of \mathbf{K}_p , providing a balanced measure that considers both the overall spread of functional types and the evenness of their contributions to the community, influenced by both their trait similarities and relative abundances. For the primary analyses in this paper, we primarily utilize $q = 1$, for which pVS simplifies to:

$$\text{pVS}(C, k, \mathbf{p}, 1) := \text{pVS}(C, k, \mathbf{p}) = \exp\left(-\sum_{i=1}^N \lambda_i \log \lambda_i\right). \quad (3)$$

However, richer insights into a community’s functional structure can be gained by considering the full spectrum of q values. Indeed, two communities can achieve identical pVS values at $q = 1$ yet achieve different pVS at other values of q (see Fig. 3). We therefore also demonstrate the flexibility of pVS by reporting *functional diversity profiles*, which are curves representing pVS as a function of q . These profiles reveal how different facets of functional diversity (richness, divergence, or evenness) are emphasized at varying sensitivities to trait variations and abundance distributions. Functional divergence is captured by pVS across all q values since pVS uses a trait similarity kernel. The q parameter then modulates how this divergence is weighted alongside functional richness and evenness. For instance, at $q = 0$, the profiles primarily reflect functional richness, highlighting the number of distinct functional types. As q increases towards $q = 1$, the profiles increasingly incorporate functional evenness, reflecting the regularity of abundance distribution across functional types. At $q > 1$, the profiles become more sensitive to the evenness and divergence of the most dominant functional types.

To better understand pVS’s behavior for the specific value of $q = 1$ when all three facets are accounted for, we consider its response under extreme conditions of trait similarity and abundance distribution.

Identical functional traits. If all members in the community possess identical functional traits, then the similarity matrix \mathbf{K} will be a matrix of all ones. In this scenario, the probability-weighted kernel matrix \mathbf{K}_p is a rank-1 matrix, with its only non-zero eigenvalue being 1. Consequently, pVS will be minimally 1. This correctly reflects a community with no functional diversity, regardless of the number of individuals or their abundance distribution.

Maximally distinct functional traits. If all members have maximally distinct functional traits (i.e., they are as dissimilar as possible), the off-diagonal elements of \mathbf{K}_p will be zero. In this scenario, the eigenvalues of \mathbf{K}_p will be equal to the relative abundances \mathbf{p} . The pVS is then simply the Hill number of order q of the community (Hill, 1973). If abundances are uniform, pVS will be N , reflecting the maximum possible functional diversity for N distinct individuals.

Uniform abundances. When all members of the community have uniform relative abundances ($p_i = \frac{1}{N}$ for all i), pVS simplifies to the Vendi Score (Friedman and Dieng, 2023; Pasarkar and Dieng, 2023). Its value is then primarily driven by the functional dissimilarities among individuals, reflecting both functional richness and divergence, while inherently demonstrating high functional evenness due to the uniform abundances.

Dominant species . If one community member is highly dominant (e.g., $p_i \approx 1$ and other $p_j \approx 0$ for $j \neq i$) then the \mathbf{K}_p matrix will be dominated by the contribution of that single species. Consequently, most eigenvalues will be very small, and pVS will approach 1, correctly indicating low functional diversity due to the overwhelming influence of a single functional type, irrespective of the total number of species or their potential functional distinctiveness.

We next highlight how pVS satisfies several essential theoretical criteria discussed in the functional diversity literature (Solow and Polasky, 1994; Mason et al., 2003; Ricotta, 2005a). All proofs are detailed in the appendix.

Theorem 2.1 (Invariance to functional redundancy).

1. *Twinning.* Consider two communities $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ with relative abundance \mathbf{p}_A and $\mathcal{C}_B = (\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{x}_{N+1})$ with relative abundance \mathbf{p}_B , where $\exists i$ such that $k(\mathbf{x}_i, \mathbf{x}_{N+1}) = 1$. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) \geq pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.
2. *Unaffected by the number of individuals/species.* Consider a community $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ with abundances $\mathbf{s}_A = (s_1, \dots, s_N)$. Consider another community $\mathcal{C}_B = (\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{x}_N, \mathbf{x}_{N+1}, \dots, \mathbf{x}_{K(N+1)})$ for some K and such that $\mathbf{t}_{N+kN+i} = \mathbf{t}_i$ and $s_{k(N+1)+i} = s_i$. Define $\mathbf{p}_A = (p_1, \dots, p_N)$ where $p_i = \frac{s_i}{\sum_{i=1}^N s_i}$. Define $\mathbf{p}_B = (\tilde{p}_1, \dots, \tilde{p}_{K(N+1)})$ where $\tilde{p}_j = \frac{s_j}{\sum_{j=1}^{K(N+1)} s_j}$. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) = pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.
3. *Invariance to species split.* Consider a community $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ with relative abundance \mathbf{p}_A . Consider another community $\mathcal{C}_B = (\mathbf{x}_1, \dots, \mathbf{x}_{N-1}, \tilde{\mathbf{x}}_N, \tilde{\mathbf{x}}_{N+1})$ such that $\mathbf{t}_N = \tilde{\mathbf{t}}_N = \tilde{\mathbf{t}}_{N+1}$ and $\mathbf{p}_B = (\mathbf{p}_{A,1}, \dots, \mathbf{p}_{A,N-1}, \tilde{\mathbf{p}}_N, \tilde{\mathbf{p}}_{N+1})$ with $\mathbf{p}_{A,N} = \tilde{\mathbf{p}}_N + \tilde{\mathbf{p}}_{N+1}$. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) = pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.

This theorem underscores pVS's robustness as an FD metric by ensuring its invariance to various forms of functional redundancy. It guarantees that the index will not be inflated by the presence of functionally identical individuals or species (twinning), by an arbitrary increase in the number of taxa that do not introduce novel functional roles, or by arbitrary taxonomic decisions (species split) that do not alter the underlying functional composition. pVS consistently focuses on the unique functional roles present in the community, rather than simply counting taxa.

Theorem 2.2 (Response to Community Composition Changes).

1. *Set monotonicity.* Consider two communities $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ and $\mathcal{C}_B = (\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{x}_{N+1})$ with relative abundances \mathbf{p}_A and \mathbf{p}_B , respectively. Assume $k(\mathbf{x}_{N+1}, \mathbf{x}_i) > 0 \ \forall i$. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) \leq pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.
2. *Distance monotonicity.* Consider two communities $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ and $\mathcal{C}_B = (\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N)$ with relative abundances \mathbf{p}_A and \mathbf{p}_B , respectively, such that $\mathbf{p}_A = \mathbf{p}_B$. Assume that $\forall i, j \in \{1, \dots, N\}$ we have $k(\mathbf{x}_i, \mathbf{x}_j) \geq k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) \geq pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.
3. *Sensitivity to functionally-rare species.* Consider a community $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$. Assume $\exists i$ such that $k(\mathbf{x}_i, \mathbf{x}_j) \geq k(\mathbf{x}_j, \mathbf{x}_k)$ and $k(\mathbf{x}_i, \mathbf{x}_k) \geq k(\mathbf{x}_j, \mathbf{x}_k) \ \forall j \neq i$ and $k \neq i$. Denote by m_i the abundance of \mathbf{x}_i . Consider the community

$\mathcal{C}_B = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ such that the abundance of \mathbf{x}_i is $m'_i \leq m_i$. Denote the relative abundances by $\mathbf{p}_A, \mathbf{p}_B$, respectively. Then $pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) \geq pVS(\mathcal{C}_B, k, \mathbf{p}_B, q)$.

4. *Continuity in abundance.* Consider a community $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$. Denote by w_i the relative abundance of \mathbf{x}_i for some i . Consider the community $\mathcal{C}_B = \mathcal{C}_A \setminus \{\mathbf{x}_i\}$. Then $\lim_{w_i \rightarrow 0} |pVS(\mathcal{C}_A, k, q) - pVS(\mathcal{C}_B, k, q)| = 0$.

This theorem highlights how pVS dynamically responds to changes in community composition, directly reflecting the facets of functional richness, divergence, and evenness. Set Monotonicity ensures that the addition of novel functional roles (increasing functional richness) is appropriately captured as an increase in diversity. Distance Monotonicity confirms that greater functional dissimilarity among individuals or species (indicating higher functional divergence) translates to higher pVS values. Lastly, the last two properties related to abundance stated in the theorem demonstrate pVS's ability to account for functional evenness by weighing the contributions of individuals (or species) by their abundance, ensuring that dominant functional types have a greater impact and that the loss of rare, unique functions is reflected in the overall diversity score. The continuity property ensures that these responses are smooth and ecologically realistic.

Theorem 2.3 (Fundamental Mathematical Properties).

1. *Boundedness and positivity.* Consider a community $\mathcal{C}_A = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ with relative abundance \mathbf{p}_A . Then $1 \leq pVS(\mathcal{C}_A, k, \mathbf{p}_A, q) \leq N$.
2. *Concavity.* Consider M non-overlapping communities $\mathcal{C}_1, \dots, \mathcal{C}_M$ such that for any $\mathbf{x} \in \mathcal{C}_m$ and $\mathbf{y} \in \mathcal{C}_n$ we have $k(\mathbf{x}, \mathbf{y}) = 0$ for all $n \neq m \in \{1, \dots, M\}$. Denote by S_m the size of \mathcal{C}_m and define $w_m = \frac{S_m}{\sum_{m'=1}^M S_{m'}}$. Denote by \mathbf{p} the relative abundance of the union of the communities and by \mathbf{p}_m the relative abundance of \mathcal{C}_m for all m . Then $pVS\left(\bigcup_{m=1}^M \mathcal{C}_m, k, \mathbf{p}, q\right) \geq \sum_{m=1}^M w_m \cdot pVS(\mathcal{C}_m, k, \mathbf{p}_m, q)$.

This theorem establishes some fundamental mathematical properties of pVS. Its positivity and boundedness between 1 and N highlight its interpretability, allowing pVS to be directly interpreted as an "effective number" of functional types. This facilitates straightforward comparisons across diverse communities, regardless of their raw richness. The concavity property is crucial for theoretical robustness. It implies that diversifying abundances (making them more even) will increase functional diversity in a predictable, diminishing-returns manner.

Collectively, the properties stated in the theorems ensure that pVS is a well-behaved, interpretable, and theoretically sound metric that can reliably quantify all three facets of functional diversity: richness (through its effective number interpretation and response to distinct additions), divergence (through its sensitivity to dissimilarity), and evenness (through its response to abundance distributions and concavity).

3 Results

To rigorously evaluate pVS and other FD indices, we utilized the simulated scenarios detailed in [Schleuter et al. \(2010\)](#). We illustrate these scenarios in Fig. 1. Only two functional traits are considered here, meaning that the functional trait space is

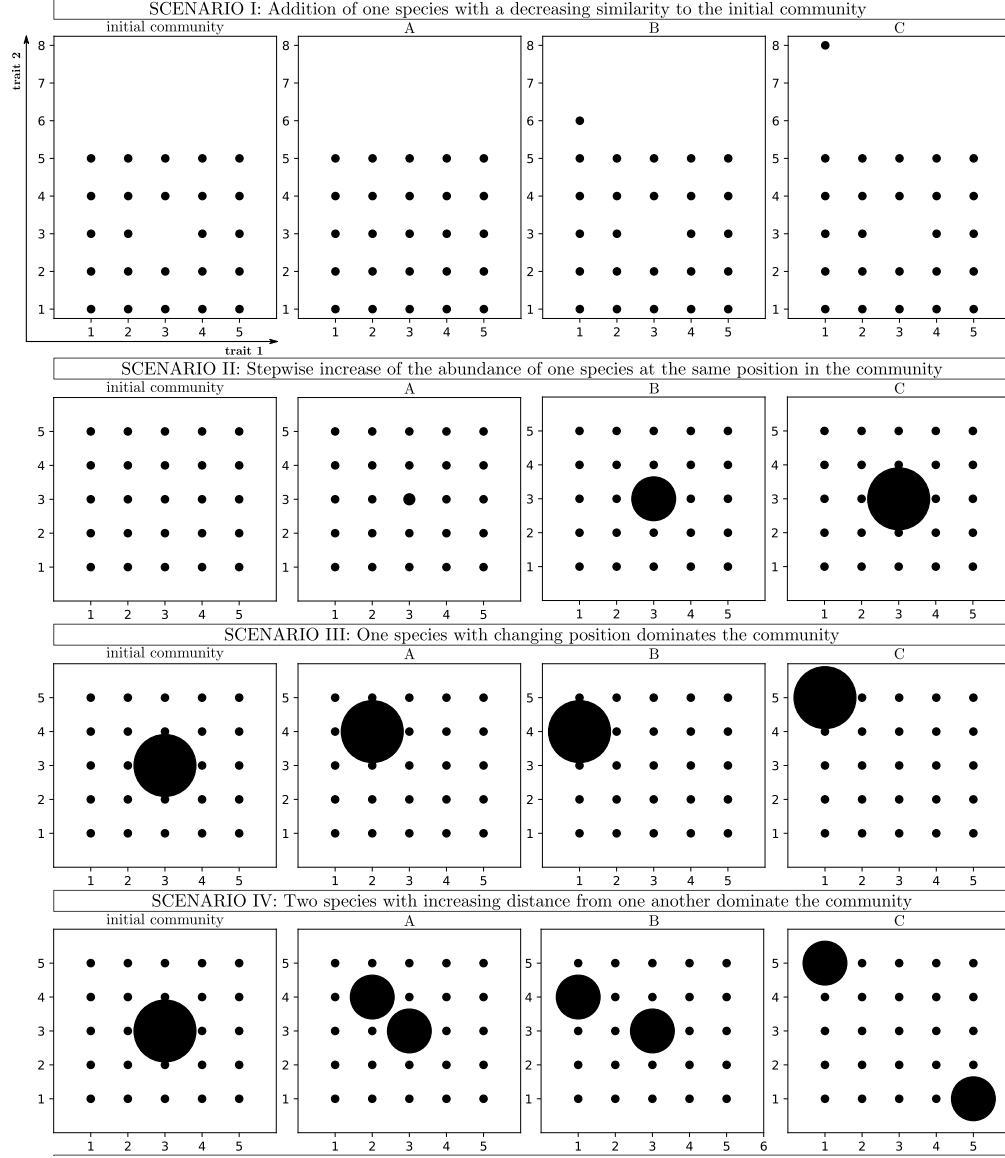


Figure 1: Simulated ecosystem scenarios with only two functional traits. A dot represents one species, and its size is correlated with its abundance.

2-dimensional. Each species is represented by a black dot, and the size of the dot is correlated with the value of its abundance: a big dot means the abundance of the species is high, and a small dot means the abundance is small. The values of the traits are integers from 1 to 8.

Scenario I corresponds to the case where one species with different trait values, and thus a different position in the functional trait space, is added to the community. Different values of similarity to the initial community are used while the abundances are kept equal for all species. Adding such a species to the community should increase functional diversity. We expect a higher increase when the similarity to the community is smaller, i.e., when the trait distance is higher.

Functional diversity index	Validity Score
Functional Attribute Diversity FaD, [Walker et al. (1999)]	1/4
modified Functional Attribute Diversity mFaD, [Schmera et al. (2009)]	1/4
Functional Diversity FD, [Petchey and Gaston (2002)]	1/4
Functional Richness, [Villéger et al. (2008)]	1/4
Average contribution, [Ricotta (2005b)]	1/4
Rao's quadratic entropy, [Botta-Dukát (2005)]	4/4
Functional Divergence FDiv, [Villéger et al. (2008)]	3/4
Functional Dispersion FDis, [Laliberté and Legendre (2010)]	4/4
Index of Mason FDVar, [Mason et al. (2003)]	1/4
FRO index to multiple traits FEve, [Villéger et al. (2008)]	1/4
Trait Onion Peeling TOP, [Fontana et al. (2015)]	1/4
Functional diversity index with NJ trees, Richness, [Cardoso et al. (2024)]	1/4
Functional diversity index with NJ trees, Dispersion, [Cardoso et al. (2024)]	3/4
Functional diversity index with NJ trees, Evenness, [Cardoso et al. (2024)]	3/4
pVS, [Friedman and Dieng (2023)]	4/4

Table 1: Performance of different FD indices on the four simulated ecosystem scenarios. The validity score is computed by recording the number of scenarios in which a given index succeeds and dividing that number by 4, the total number of scenarios. Most metrics fail in at least 1 scenario. The pVS, Rao's quadratic entropy, and FDis are the only metrics that succeed in all scenarios.

Scenario II consists in increasing the abundance of only one species in the community. We start with a balanced community (all abundances are equal), and we give more and more importance to one species. We expect a decrease in functional diversity as one species dominates in abundance. Indeed, an assemblage of 10 species dominated by one is less diverse than another with 10 equally represented species.

Scenario III corresponds to the case where one species dominates the community in terms of abundance, but we now vary its position in the trait space. The similarity of the dominant species compared to the community varies. We expect an increase in functional diversity when the similarity decreases.

In Scenario IV, two species dominate the community and have the same abundance. We increase the distance between these two species. The less similar they are, the more functionally diverse the community is supposed to be.

Trait distance is computed using the L_1 distance. pVS uses a Radial Basis Function kernel to define similarity from these distances. The similarity kernel is

$$k(x, y) = \exp\left(\frac{-d(x, y)^2}{2\sigma^2}\right) \quad (4)$$

where $d(x, y)$ is the distance between x and y and σ^2 is the width of k .

To quantify performance, we compute a *validity score* for each index. We define the validity score as the number of scenarios in which the index exhibited the

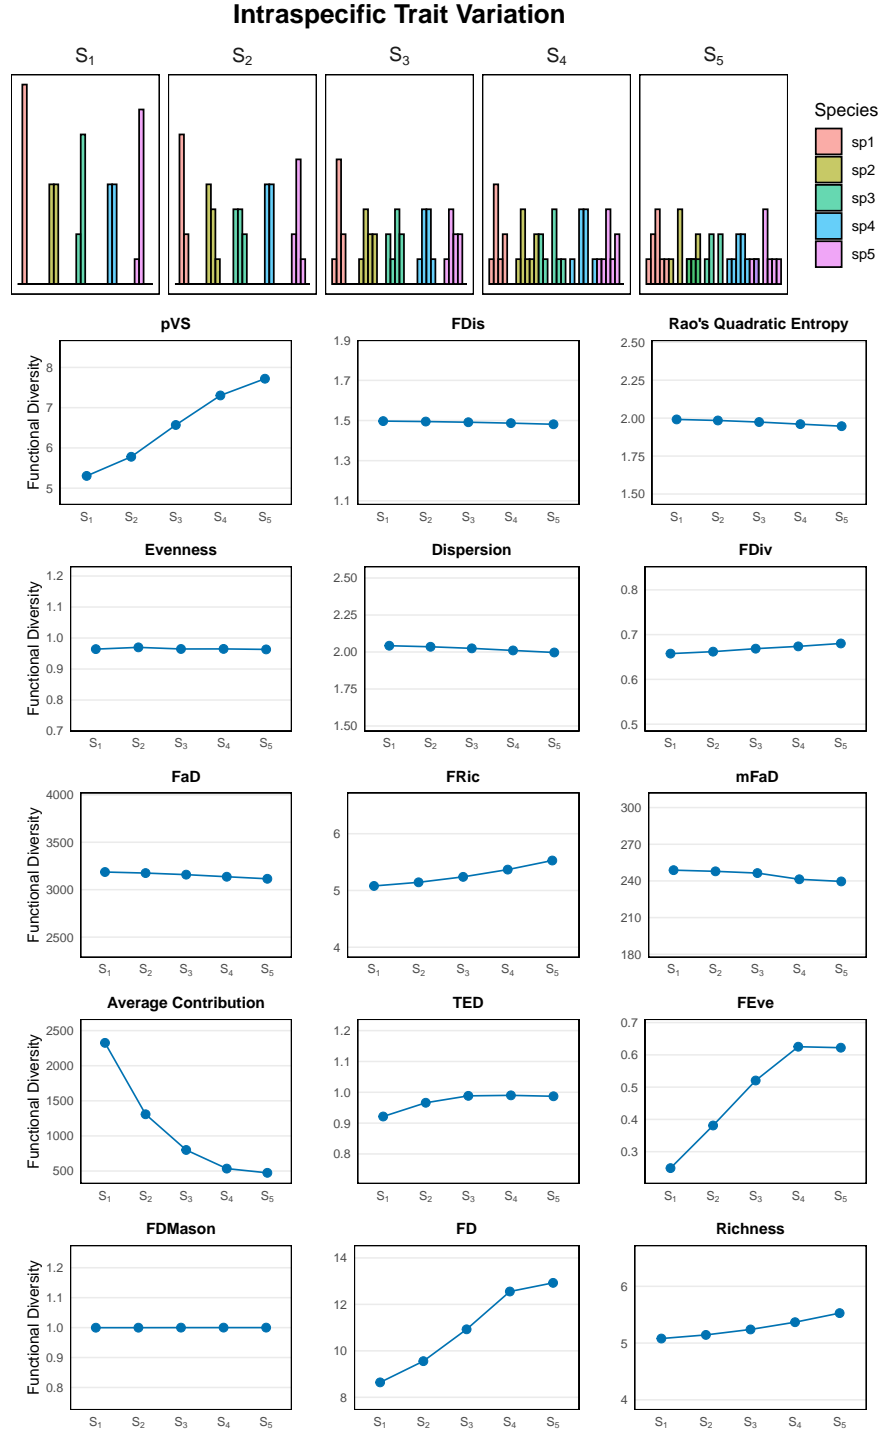


Figure 2: Top: Histograms of trait values for synthetic communities with increasing levels of variation from left to right. **Bottom:** Behavior of 15 different FD indices, computed at the individual level, as intraspecific trait variance increases. Only pVS, FDiv, FRic, FDMason, FD, and Richness increase monotonically across the 5 communities. FDMason and FDiv both increase by less than 0.1% and 4% between community S_1 and community S_5 despite the drastic changes to the community's traits. Rao's quadratic entropy and FDis fail to capture ITV. pVS is the only metric that is both accurately sensitive to ITV and passes all scenarios from Table 1.

expected ecological response to community changes divided by the total number of scenarios. As shown in Table 1, 9 out of the 15 FD indices only succeed in 1 scenario and fail on all the others. On the other hand, 3 out of the 15 metrics fail on 1 scenario, with Rao's quadratic entropy [Ricotta \(2005a\)](#), Functional Dispersion (FDis) ([Laliberté and Legendre, 2010](#)), and pVS being the only indices that succeed in all scenarios.

We additionally show that pVS can account for intraspecific trait variation (ITV), unlike the metrics that pass the majority of scenarios in Table 1, i.e. Rao's quadratic entropy and Functional Dispersion (FDis). In Figure 2, we simulate communities of 5 species with 1 trait, and increasing the levels of trait variance for each species. Intuitively, we expect functional diversity to increase when we introduce more intraspecific variance, yet most metrics are either constant or decreasing. Among those that are sensitive to ITV, pVS is the only one to pass all scenarios in Table 1. (Note: we omit the TOP metric here because it requires at least two-dimensional trait data.)

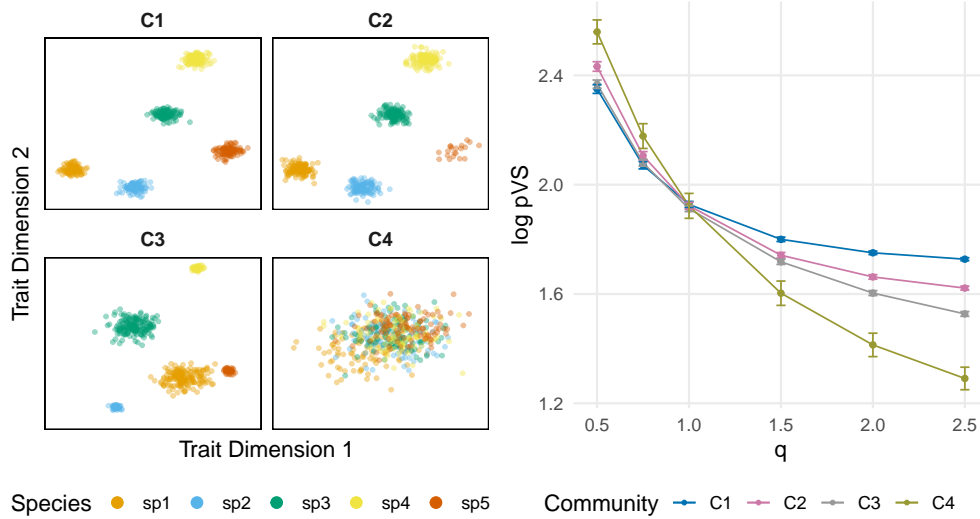


Figure 3: The order q of pVS prioritizes distinct community characteristics. **Left:** 4 communities consisting of 5 species with 2 traits. C1 contains 5 species of equal size. C2 has 1 rare species and slightly higher variance within the other 4 species. C3 has two dominant species. C4 has 5 similar species that all have high variance. **Right:** the **FD profile** of the four different communities as described by the curve of \log -pVS as a function of q across 10 simulations. All communities have almost the same pVS at $q = 1$ (difference within 1 standard deviation), but have pVS that differ significantly at other values of q . Community C4 has the highest pVS at $q=0.5$, but the lowest at larger values of q . Community C1 has the lowest pVS at $q=0.5$, but the highest at larger values of q .

Finally, we empirically show how the order q used in the pVS will prioritize different aspects of functional diversity. In Figure 3, we highlight 4 examples of these differences. When $q < 1$, the pVS places greater emphasis on the presence of rare species and the overall spread of trait values, leading to higher pVS for communities with broader overall trait distributions, like communities C2 and C4. As q increases,

the pVS becomes increasingly sensitive to the size of dominant species, which explains why communities with more even abundances across species have higher pVS, like community C1. In contrast, communities with many highly similar trait values, like community C4, tend to drop in rank. Community C3 has two distinct dominant species, giving it a higher pVS than C4 for large q . Tuning q allows researchers to focus on either rarity and dispersion (low q) or dominance and functional redundancy (high q), making pVS a flexible tool for studying various ecological communities.

4 Conclusion

The robust and comprehensive quantification of functional diversity (FD) is a long-standing challenge in ecology. Existing metrics struggle to simultaneously capture richness, divergence, and evenness, and often neglect intraspecific trait variation (ITV). In this paper, we demonstrate that the probability-weighted Vendi Score (pVS) can be adapted to functional ecology to provide a robust and widely-applicable FD index. Defined as the exponential of the Rényi entropy of the eigenvalues of the abundance-weighted trait similarity matrix, pVS can be applied at the individual level, which naturally accounts for ITV and provides a more fine-grained FD assessment. It is also readily applicable at the species level by aggregating individual data. We rigorously prove that pVS satisfies several essential theoretical criteria for robust FD metrics and consistently exhibits the expected ground-truth behavior across simulated ecosystem scenarios where many existing functional diversity metrics fail. By unifying the multifaceted nature of FD and incorporating ITV within a single, theoretically sound framework, pVS offers a powerful, interpretable, and comprehensive tool for ecological research.

Code and Data Availability

The Python code to compute the probability-weighted Vendi Score (pVS) on your data is freely available in the Vertaix GitHub at <https://github.com/vertaix/Vendi-Score>. The pVS is also available as a Python package at <https://pypi.org/project/vendi-score/>. The pVS is also readily available in R in the same GitHub at <https://github.com/vertaix/Vendi-Score-R>.

All the results in this paper can be reproduced using the R code in the repo <https://github.com/vertaix/pVS-Functional-Diversity>. That same repo contains the code to generate the data for the ecosystem scenarios and reproduce Table 1.

References

- Botta-Dukát, Z. (2005). Rao's quadratic entropy as a measure of functional diversity based on multiple traits. *Journal of Vegetation Science*, 16(5):533–540.
- Cardoso, P., Guillerme, T., Mammola, S., Matthews, T. J., Rigal, F., Graco-Roza, C., Stahls, G., and Carlos Carvalho, J. (2024). Calculating functional diversity metrics using neighbor-joining trees. *Ecography*, 2024(7).

- D’Orazio, M. (2021). Distances with mixed type variables some modified gower’s coefficients.
- Fontana, S., Petchey, O. L., and Pomati, F. (2015). Individual-level trait diversity concepts and indices to comprehensively describe community change in multidimensional trait space. *Functional Ecology*, 30(5):808–818.
- Friedman, D. and Dieng, A. B. (2023). The vendi score: A diversity evaluation metric for machine learning. *arXiv preprint arXiv:2210.02410*.
- Hill, M. O. (1973). Diversity and evenness: a unifying notation and its consequences. *Ecology*, 54(2):427–432.
- Laliberté, E. and Legendre, P. (2010). A distance-based framework for measuring functional diversity from multiple traits. *Ecology*, 91(1):299–305.
- Mason, N. W., MacGillivray, K., Steel, J. B., and Wilson, J. B. (2003). An index of functional diversity. *Journal of Vegetation Science*, 14(4):571–578.
- Mouchet, M. A., Villéger, S., Mason, N. W. H., and Mouillot, D. (2010). Functional diversity measures: an overview of their redundancy and their ability to discriminate community assembly rules. *Functional Ecology*, 24(4):867–876.
- Palacio, F. X., Ottaviani, G., Mammola, S., Graco-Roza, C., de Bello, F., and Carmona, C. (2024). Integrating intraspecific trait variability in functional diversity: An overview of methods and a guide for ecologists.
- Pasarkar, A. P. and Dieng, A. B. (2023). Cousins of the vendi score: A family of similarity-based diversity metrics for science and machine learning. *arXiv preprint arXiv:2310.12952*.
- Pavoine, S., Vallet, J., Dufour, A., Gachet, S., and Daniel, H. (2009). On the challenge of treating various types of variables: application for improving the measurement of functional diversity. *Oikos*, 118(3):391–402.
- Petchey, O. L. and Gaston, K. J. (2002). Functional diversity (fd), species richness and community composition. *Ecology Letters*, 5(3):402–411.
- Ricotta, C. (2005a). A note on functional diversity measures. *Basic and Applied Ecology*, 6(5):479–486.
- Ricotta, C. (2005b). A note on functional diversity measures. *Basic and applied Ecology*, 6(5):479–486.
- Schleuter, D., Daufresne, M., Massol, F., and Argillier, C. (2010). A user’s guide to functional diversity indices. *Ecological Monographs*, 80(3):469–484.
- Schmera, D., Podani, J., and Erős, T. (2009). Measuring the contribution of community members to functional diversity. *Oikos*, 118(7):961–971.
- Schmera, D., Ricotta, C., and Podani, J. (2023). Components of functional diversity revisited: A new classification and its theoretical and practical implications. *Ecology and Evolution*, 13(10).
- Solow, A. R. and Polasky, S. (1994). Measuring biological diversity. *Environmental and Ecological Statistics*, 1(2):95–103.

- Villéger, S., Mason, N. W. H., and Mouillot, D. (2008). New multidimensional functional diversity indices for a multifaceted framework in functional ecology. *Ecology*, 89(8):2290–2301.
- Walker, B., Kinzig, A., and Langridge, J. (1999). Plant attribute diversity, resilience, and ecosystem function: the nature and significance of dominant and minor species. *Ecosystems*, 2(2):95–113.