

Soundscape perception indices (SPIs): Developing context-dependent single value scores of multidimensional soundscape perceptual quality^{a)}

Andrew Mitchell,^{1,2,b)}  Francesco Aletta,²  Tin Oberman,²  and Jian Kang² 

¹Bartlett School of Sustainable Construction, University College London, London, United Kingdom

²Institute for Environmental Design and Engineering, University College London, London, United Kingdom

ABSTRACT:

The soundscape approach provides a basis for considering the holistic perception of sound environments in context. Whereas steady advancements have been made in methods for assessment and analysis, a gap exists for comparing soundscapes and quantifying improvements in the multidimensional perception of a soundscape. To this end, there is a need for the creation of single value indices to compare soundscape quality which incorporate context, aural diversity, and specific design goals for a given application. Just as a variety of decibel-based indices have been developed for various purposes (e.g., L_{Aeq} , L_{Ceq} , L_{90} , L_{den} , etc.), the soundscape approach requires the ability to create original indices for different uses, which share a common language and understanding. Therefore, a unified framework for creating bespoke and reference single index measures of soundscape perception is proposed, allowing for different metrics to be defined in the future. This framework is based on a four-step test-target paradigm wherein a desired soundscape perception is defined as a target distribution within the soundscape circumplex, and the two-dimensional Kolmogorov-Smirnov distance is used to test an assessed soundscape against this target. Applications and implications of this framework are discussed, and a multi-objective optimisation method for empirically defining perception indices is proposed. © 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.1121/10.0034417>

(Received 27 June 2024; revised 25 October 2024; accepted 26 October 2024; published online 3 December 2024)

[Editor: James F. Lynch]

Pages: 3694–3706

I. INTRODUCTION

The European Union (EU) Green Paper on Future Noise Policy indicates that 80×10^6 EU citizens are suffering from potentially harmful environmental noise levels, according to the World Health Organization (WHO) recommendations (Berglund *et al.*, 1999). The publication of the EU Directive Relating to the Assessment and Management of Environmental Noise (END) (European Union, 2002) more than two decades ago has led to major actions across Europe, where reducing noise levels are their main focus, for which billions of Euros are being spent. However, it is widely recognized that solely reducing sound level in people's living environments is not always feasible or cost-effective and, more importantly, with only ~30% of environmental noise annoyance depending on physical aspects of the signal, such as acoustic energy (Guski, 1997), sound level reduction will not necessarily lead to improved quality of life. For this reason, from a public health point of view, it is necessary to explore alternative management and design strategies for acoustic environments that rely on more positive soundscapes rather than merely environments not

affected by noise pollution (Aletta *et al.*, 2018; Kang, 2023; Kang *et al.*, 2023).

Soundscape design, separate from (and complementary to) noise control engineering, is about the relationships between human physiology, perception, the sound environment, and its sociocultural context (Kang, 2006). Soundscape research represents a paradigm shift in that it combines physical, social, and psychological approaches and considers environmental sounds as a “resource” rather than “waste” (Kang and Schulte-Fortkamp, 2016), relating to perceptual constructs rather than just physical phenomena. However, the current research is still at the stage of describing and identifying the problems and tends to be fragmented and focussed on only special cases, e.g., subjective evaluations of soundscapes for residential areas (Chen and Kang, 2023; Schulte-Fortkamp and Kang, 2013). In the movement from noise control to soundscape creation (Aletta and Kang, 2015), a vital step is the standardization of methods to assess soundscape quality.

A common aim for implementing soundscape assessment in practice is to compare the quality of different soundscapes. Often (but not always) the goal is to identify a “good” soundscape compared to a “bad” soundscape. However, this presents several challenges:

^{a)}This paper is part of a special issue on Advances in Soundscape: Emerging Trends and Challenges in Research and Practice.

^{b)}Email: a.j.mitchell@ucl.ac.uk

- What makes a soundscape good or bad is highly contextual; that is, the same acoustic environment may result in different appreciations and perceptual outcomes depending on where/when it is happening and what groups of individuals are there to experience it.
- On what metric should the quality rating be based? Previous attempts at defining objective metrics of “soundscape quality” assessments have fallen short of capturing the multidimensionality of people’s perception of surrounding acoustic environments.
- How can we deal with different requirements and definitions of how a soundscape should be perceived? Soundscape constructs are normally seen as highly individualised, whereas designing the soundscapes of public spaces should look at accommodating the needs of a given community of a space as a whole.

In many cases, the ultimate aim is to be able to rank soundscapes based on their quality. There is pressure from stakeholders and policymakers to move toward such simplified assessment protocols. However, any ranking metric should be flexible and able to handle a variety of contexts and definitions of what a good soundscape is for a given purpose. To address this, we will propose the Soundscape Perception Index (SPI) framework, a flexible method for defining single value indices of soundscape quality based on distributions within the soundscape circumplex model (SCM; *Axelsson et al., 2010, 2012; Mitchell et al., 2022*). As previously suggested, the primary motivation behind the development of the SPI framework stems from the need to address the existing gap in quantifying and comparing soundscape quality across diverse contexts and applications. By creating a unified framework for defining these indices, the aim is to empower stakeholders, decision-makers, and researchers with the ability to create tailored indices that align with their specific objectives and design goals while simultaneously enabling cross-comparisons and benchmarking against empirically defined reference soundscapes. This dual approach not only acknowledges the context-dependent nature of soundscape perception but also fosters a common language and understanding, facilitating knowledge sharing and collaborative efforts within the field. This paper will demonstrate the SPI framework and test whether it is capable of scoring soundscape quality and generating consistent rankings of soundscapes across different contexts.

II. THEORETICAL BACKGROUND

In *Aletta et al. (2016)*, the authors defined a framework for categorising the components of a soundscape assessment. They define three aspects: soundscape descriptors, soundscape indicators, and soundscape indices. Soundscape descriptors are defined as “measures of how people perceive the acoustic environment,” and soundscape indicators are defined as “measures used to predict the value of a soundscape descriptor.” The relationship between soundscape indicator(s) and a soundscape descriptor effectively defines what has been previously referred to as a “predictive

soundscape model” (*Aletta et al., 2016; Mitchell, 2022*). There are primarily two rationales for modelling the relationship between the physical attributes and perceived (i.e., soundscape) qualities of the acoustic environment. First, a predictive model can forecast how individuals would perceive the acoustic environment, eliminating the need for labour-intensive surveys (*Mitchell et al., 2023*). Second, a precise predictive model may unveil the root causes of these perceived qualities, thereby serving as a valuable tool for design. *Lionello et al. (2020)* provided a review of such models and concluded that contextual features play an important role in increasing the quality of the model. Indices, on the other hand, the primary focus of this article, are single numerical values that combine multiple indicators or descriptors to provide a comprehensive representation of the overall soundscape perception and allow for comparison between soundscapes.

The earliest and most commonly used scientific index measuring sound level is the decibel (dB). To represent the overall level of sound with a single value on one scale, as the dB index does, is often desirable. For this purpose, a number of different values representing sounds at various frequencies must be combined. Several frequency weighting networks have been developed since the 1930s, considering typical human responses to sound based on equal-loudness-level contours (*Fletcher and Munson, 1933*) and, among them, the A-weighting network, with resultant dB values called dBA, has been commonly used in almost all of the national/international regulations (*Kryter, 1994*). However, there have been numerous criticisms on its effectiveness (*Parmanen, 2007*) as the correlations between dBA and perceived sound quality (e.g., noise annoyance) are often low (*Hellman and Zwicker, 1987*).

Another set of indices is psychoacoustic magnitudes, including loudness, fluctuation strength, or roughness, sharpness, and pitch strength, developed through sound quality studies of industrial products since the 1980s (*Zwicker and Fastl, 2007*). These emerged when it was conceived that acoustic emissions can be characterised beyond just sound level (*Blauert and Jekosch, 1997*). However, although psychoacoustic magnitudes have proven to be successful for the assessment of product sound quality, in the field of environmental acoustics, their applicability has been limited (*Fastl, 2006*) because a significant feature of environmental acoustics is that there are multiple/dynamic sound sources. Additionally, whereas psychoacoustic magnitudes incorporate perceptual aspects, both dB-based and psychoacoustic indices are ultimately describing the acoustic signal and not the soundscape perception and may, therefore, be more accurately described as indicators rather than soundscape indices (*Mitchell et al., 2023*).

When applied to urban sound and, specifically, noise pollution, the soundscape approach introduces three key considerations beyond traditional noise control methods:

- (1) Considering all aspects of the environment which may influence perception and not just the sound level and

spectral content (e.g., visual setting, odour environment, spatial layout, etc.);

- (2) an increased and integrated consideration of the varying impacts which different sound sources and sonic characteristics have on perception; and
- (3) a consideration of the positive and negative dimensions of soundscape perception.

This approach can enable better outcomes by identifying positive soundscapes (in line with the END's mandate to "preserve environmental noise quality where it is good"; European Union, 2002), better identifying specific sources of noise which impact soundscape quality and pinpointing the characteristics that may need to be decreased, and illuminating alternative methods which could be introduced to improve a soundscape in which a reduction of noise is impractical (Fiebig, 2018; Kang and Aletta, 2018). Factors such as the presence of natural or human-made sounds, their temporal patterns, and the overall contextual meaning ascribed to these sounds all contribute to the holistic perception of a soundscape.

A. Existing "soundscape indices"

Although the field of soundscape research has witnessed substantial progress, the development of standardized indices for evaluating and comparing soundscapes across diverse contexts has been relatively limited. Existing indices can be broadly observed as arising from two domains: soundscape ecology and soundscape perception.

1. Soundscape ecology and bioacoustics

Within the realm of soundscape ecology, indices such as the acoustic diversity index (ADI) and frequency-dependent acoustic diversity index (FADI; Xu *et al.*, 2023) have been developed to quantify the diversity and complexity of acoustic signals within a given soundscape. Similar indices [e.g., ADI, normalised difference soundscape index (NDSI), acoustic complexity index (ACI)] have also been developed to analyse the acoustic signal of complex acoustic environments and indicate the richness and diversity of biophonic (natural) and anthrophonic (human-made) sound sources. However, whereas these indices contribute valuable insights into the ecological aspects of soundscapes, they do not directly address the perceptual dimensions that are central to the soundscape approach (Schulte-Fortkamp *et al.*, 2023). The multidimensional nature of soundscape perception, encompassing factors such as pleasantness, eventfulness, and familiarity, necessitates a more comprehensive and context-sensitive approach.

2. Soundscape perception

In the domain of soundscape perception, several indices have emerged as attempts to quantify the perceived quality of soundscapes, particularly, in urban environments.

The green soundscape index (GSI; Kogan *et al.*, 2018) incorporates factors such as the presence and levels of natural sounds and human-made sounds and their respective contributions to the overall soundscape perception. The GSI

is defined as the ratio of the perceived extent of natural sounds (PNS) to the perceived extent of traffic noise (PTN), ranging between one-fifth and five. Subsequently, Cao and colleagues (Cao *et al.*, 2020; Yang *et al.*, 2022) proposed the red soundscape index (RSI), defined as the ratio of the perceived extent of human sounds to the perceived extent of either natural sounds (RSI_n) or traffic sounds (RSI_t), arguing that the GSI alone was not suitable for all urban soundscape design applications

Building on ecological diversity concepts, Liu *et al.* (2014) introduced the soundscape diversity index (SDI), which quantifies the probability of two randomly selected sounds in a soundscape belonging to different categories, providing a measure of soundscape complexity. Expanding on this approach, Xiang *et al.* (2023) defined an expanded set of soundscape diversity indices, including the SDI, the soundscape richness index (SRI), the soundscape dominance index (SDO), and the soundscape evenness index (SEI). These indices, adapted from species diversity measures in ecology, offer a more nuanced approach to quantifying aspects of soundscape perception. Xiang *et al.* (2023) demonstrated that these indices could be partially explained by existing acoustic indicators and were more suitable for evaluating urban green spaces than traditional acoustic indices.

Guo *et al.* (2023) proposed the harmonious degree of sound sources (SHD) index, which combines perceived loudness, occurrence, and preference for sound sources. The SHD assess how well the dominance of a sound aligns with visitors' preferences, aiming to reflect the harmonious status of sounds in a soundscape.

Kang *et al.* (2019) proposed the development of a set of soundscape indices which might take the form $SSID = f(\text{physical factors}) + f(\text{contextual factors}) + \dots$, where the functions and weights of each aspect influencing soundscape perception (i.e., physical/acoustic parameters, contextual and visual factors, personal factors, etc.) could be derived statistically from a large dataset of soundscape surveys. The work presented here represents a development of this thinking which has grown out of the SSID project, where the analysis and indexing of perception data and connection between soundscape indicators and perception have been separated (Mitchell *et al.*, 2023). This modularisation of perception prediction based on objective factors and soundscape index creation should enable more sophisticated and thoughtful index creation and more advanced and updateable prediction models.

Although these indices offer valuable insight into specific aspects of soundscape perception, they are limited in their ability to capture the full multidimensionality of soundscape experience across diverse contexts. The SPI framework presented in this paper builds on these efforts by providing a flexible, context-sensitive approach to soundscape assessment. Unlike many previous indices, the SPI is not an analysis of an acoustic signal but rather is an index of perception based on soundscape descriptors. Furthermore, it does not represent a single target in a particular context but is a generalisable, extensible, and adaptable framework for scoring soundscapes against any goal defined by the user.

The remainder of the paper will introduce and demonstrate this framework, providing a case study of defining an appropriate target.

III. ESTABLISHING THE SPI FRAMEWORK

The index framework, SPI, introduced in this paper is defined here as the agreement between an observed or modelled soundscape perception distribution and a target soundscape perception distribution. Its goal is to determine whether a soundscape—be it a real-world location, a proposed design, or a hypothetical scenario—aligns with the desired (or reference) perception of that soundscape. This is achieved by first defining the target distribution, which could represent what is considered to be the “ideal” soundscape perception for a given context or application. The test distribution is then compared to the target distribution using a distance metric, which quantifies the deviation between the two distributions. The resulting distance value serves as the basis for calculating the SPI, where smaller distances indicate a closer alignment between the perceived soundscape and target soundscape perception.

We refer to this as an index framework rather than a single index as the SPI can be tailored to specific contexts and applications by defining a range of target distributions. A single index is, thus, created for each target distribution. A SPI value, therefore, does not represent a good or bad soundscape but rather a measure of how closely the perceived soundscape aligns with the desired target soundscape perception. This approach allows for the development of bespoke indices tailored to specific design goals and objectives while also enabling cross-comparisons and benchmarking against empirically defined reference soundscape targets.

A. Core method

SPI is grounded in the SCM (Axelsson *et al.*, 2010, 2012), a robust theoretical foundation for understanding and representing the multidimensional nature of soundscape perception. The reason for grounding the SPI in the soundscape circumplex is that we have observed this model to become the most prevalent assessment model in soundscape

literature (Aletta and Torresin, 2023). The SCM is built on a series of descriptors referred to as the perceived affective quality (PAQ), proposed by Axelsson *et al.* (2010). These PAQs are based on the pleasantness-eventfulness paradigm (and its original version, the valence-arousal paradigm), which has been adopted in research on emotions and environmental psychology, in particular, Russell’s circumplex model of affect (Russell, 1980). As summarised by Axelsson *et al.* (2010), “Russell’s model identifies two dimensions related to the perceived pleasantness of environments and how activating or arousing the environment is.”

One benefit of the circumplex model is that, as a whole, it encapsulates several of the other proposed soundscape descriptors—in particular, annoyance, pleasantness, tranquility, and possibly restorativeness (Aletta *et al.*, 2016). According to Axelsson (2015), the two-dimensional circumplex model of PAQ provides the most comprehensive information for soundscape assessment. It is also possible that the overall soundscape quality could itself be derived from the pleasant-eventful scores derived for a soundscape. The circumplex also lends itself well to questionnaire-based methods of data collection, as proposed in ISO/TS 12913-2 (2018). In contrast to methods such as soundwalks, interviews, and laboratory experiments, *in situ* questionnaires are able to provide the quality and amount of data which is necessary for statistical modelling. Combined, these factors make the circumplex most appropriate for a single index as it provides a comprehensive summary of soundscape perception.

There are four steps involved in calculating the SPI, as shown in Fig. 1:

- (1) Define and parameterise the target circumplex distribution;
- (2) sample the target distribution and prepare the test distribution;
- (3) compare test and target distributions using the distance metric [two-dimensional Kolmogorov-Smirnov (KS) distance, D_{BKS}]; and
- (4) calculate $SPI = 100(1 - D_{BKS})$.

These steps and their required background are discussed in detail in Secs. III B–III E. Section IV B 1 will, then,

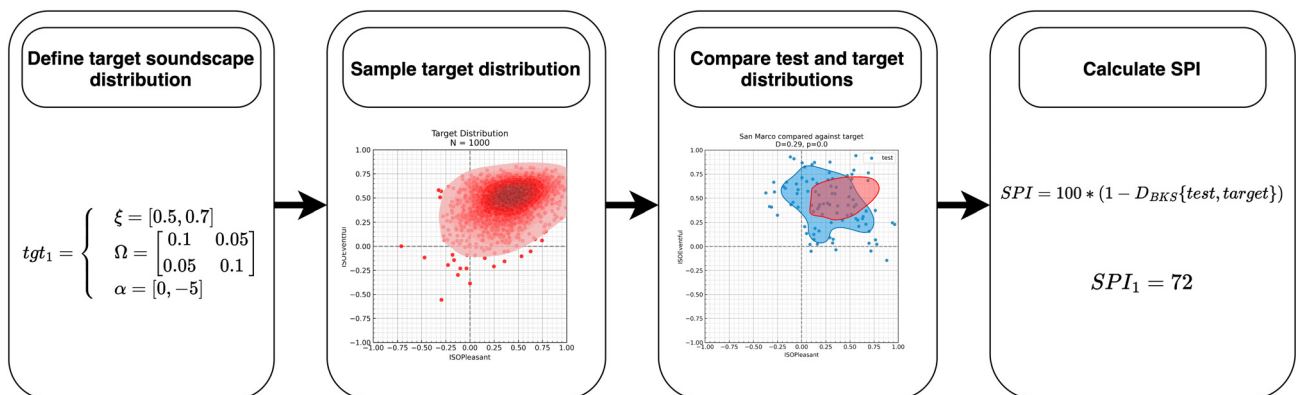


FIG. 1. (Color online) Steps for calculating the SPI.

present strategies for defining targets and their applications. Throughout this paper, we use the data contained in the International Soundscape Database (ISD; Mitchell *et al.*, 2024), which includes 1300+ individual responses on the PAQ scales, collected across 13 locations in London and Venice, according to the SSID Protocol (Mitchell *et al.*, 2020).

B. Define and parameterise a soundscape circumplex distribution

To move the eight-item PAQ responses into the two-dimensional circumplex space, we use the projection method that was first presented in ISO/TS 12913-3 (ISO/TS 12913-3, 2019). This projection method and its associated formulas were recently updated further in Aletta *et al.* (2024) to include a correction for the language in which the survey was conducted. Aletta *et al.* (2024) also provides adjusted angles for translations of the circumplex attributes to be used in calculating the P_{ISO} and E_{ISO} coordinates. Once the individual perceptual responses are projected into the circumplex space, the resulting data for each location are treated as a circumplex distribution. There are several advancements in considering circumplex distributions compared to the discussions originally given in Mitchell *et al.* (2022), which are necessary for SPI. Before exploring the SPI method and target setting more specifically, we will first address these developments.

The circumplex is defined by two axes, P_{ISO} and E_{ISO} , which are limited to the range $[-1, +1]$. Typically, data in the soundscape circumplex is treated as a combination of two independent normal distributions, one for each axis (Mitchell *et al.*, 2022; Ooi *et al.*, 2022). In some applications, this approach is sufficient for capturing the distribution of soundscape perception, however, defining a target distribution for SPI requires a more precise approach. The independent normal distribution approach relies on three key assumptions:

- (1) The two axes are normally distributed;
- (2) the two axes are symmetrically distributed; and
- (3) the two axes are independent of each other.

Whereas the first assumption is generally valid, the second and third assumptions are not always met in practice. In particular, the distribution of soundscape perception responses in the circumplex is often characterised by a high degree of skewness, which can lead to inaccuracies in the calculation of the SPI. Soundscape circumplex distributions are most appropriately described as a bivariate skew-normal distribution (Azzalini, 2005), which accurately reflects the relationship between the two dimensions of the circumplex and the fact that real-world perceptual distributions have been consistently observed to not be strictly symmetric.

The skew-normal distribution is defined by three parameters: location (μ), scale (σ), and shape (α). The location parameter defines the centre of the distribution, the scale parameter defines the spread of the distribution, and the shape parameter defines the skew of the distribution. The one-dimensional skew-normal distribution is defined as (Azzalini and Valle, 1996)

$$\phi(z; \alpha) = 2\phi(z)\Phi(\alpha z) \quad \text{for } z \in \mathbb{R}, \quad (1)$$

where ϕ and Φ are the standard normal probability density function and distribution function, respectively, and α is a shape variable which regulates the skewness. The distribution reduces to a standard normal density when $\alpha = 0$. The bivariate skew-normal distribution extends this concept to two dimensions, allowing for the modelling of asymmetric and skewed distributions in a two-dimensional space such as the soundscape circumplex. The multivariate skew-normal (MSN) distribution, including scale and location parameters, is given by combining the normal density and distribution functions (Azzalini and Capitanio, 1999) such that

$$Y = 2\phi_k(y - \xi; \Omega)\Phi\{\alpha^T \omega^{-1}(y - \xi)\}, \quad (2)$$

where ϕ_k is the k -dimensional normal density with location ξ , shape α , and covariance matrix Ω . $\Phi\{\cdot\}$ is the normal distribution function, and α is a k -dimensional shape vector. When $\alpha = 0$, Y reduces to the standard multivariate normal $N_k(\xi, \Omega)$ density. A circumplex distribution can, therefore, be parameterised¹ with a 2×2 covariance matrix Ω , a 2×1 location vector ξ , and a 2×1 shape vector α , written as

$$Y \sim \text{MSN}(\xi, \Omega, \alpha). \quad (3)$$

By fitting a MSN distribution to empirical soundscape perception responses, it becomes possible to accurately capture the asymmetry and skewness of the distribution. A bivariate skew-normal distribution can be summarised as a set of these three parameters. Once parameterised, the distribution can then be sampled from to generate a synthetic distribution of soundscape perception responses.

Soundscape targets can, thus, be set by defining the desired MSN distribution. To demonstrate this, we will construct three arbitrary targets, which will be used later to score three SPIs. The parameters chosen for the example targets are given in Table I.

C. Sample a target distribution

Once the parameters for a MSN are defined (i.e., the target), the MSN is then sampled using the `sn` package

TABLE I. The MSN direct parameterisations for three arbitrary example target distributions. `tgt1` is located in the pleasant half with a wide variance and a positive skew along the pleasantness axis. `tgt2` is located in the calm quadrant with a typical variance and a negative skew along the pleasantness axis and a positive skew along the eventful axis. `tgt3` is located in the vibrant quadrant with a moderate variance and a negative skew along the eventfulness axis.

Target	Location, ξ	Covariance matrix, Ω	Shape, α
<code>tgt₁</code>	[0.5, 0.0]	$\begin{bmatrix} 0.2 & 0.0 \\ 0.0 & 0.2 \end{bmatrix}$	[1, 0]
<code>tgt₂</code>	[1.0, -0.4]	$\begin{bmatrix} 0.18 & -0.04 \\ -0.04 & 0.09 \end{bmatrix}$	[-8, 1]
<code>tgt₃</code>	[0.5, 0.7]	$\begin{bmatrix} 0.1 & 0.05 \\ 0.05 & 0.1 \end{bmatrix}$	[0, -5]




FIG. 2. (Color online) Example of defining and sampling from three arbitrary bespoke targets.

(Azzalini, 2021) in R (R Core Team, 2018). This is to prepare the target distribution to be compared with the empirical test distribution. Several restrictions to the possible parameter values apply, most importantly the covariance matrix Ω must be a positive-definite matrix. In depth discussions of how these parameterisations should be defined and their restrictions can be found in Azzalini (2016). Figure 2 shows the result of sampling ($n = 1000$) the three example distributions given in Table I and plotting them as soundscape distributions.

D. Compare the target and test soundscape assessment distributions

Central to the SPI framework is the concept of a distance metric, which quantifies the deviation of a given soundscape from a desired target soundscape. This distance metric serves as the basis for calculating the SPI value, where smaller distances indicate a closer alignment between the perceived soundscape and target soundscape perception. The distance between the test and target soundscape distributions is calculated using a two-dimensional KS distance (D_{BKS} ; Fasano and Franceschini, 1987). The KS distance is a nonparametric metric of the equality of continuous distributions, which is sensitive to the locations and shapes of the distributions (Chakravati et al., 1967).

Essentially, we approach this as a problem of (dis)similarity between soundscapes. The D_{BKS} distance metric is then proposed to assess how similar any two given soundscapes distributions are within the circumplex. Taken to the extreme, two perfectly matching distributions in the soundscape circumplex would return a 100% SPI value, whereas two completely dissimilar distributions would return a 0% SPI value. In practical terms, for the former, this will never be achieved in real-world scenarios; for the latter, it is also difficult to estimate how low the SPI value could actually go, and it should be considered that the distance may happen in different directions within the circumplex space. For instance, if a distribution for a vibrant soundscape was taken as a reference, a compared soundscape distribution may

exhibit low SPI values for being located in the calm, or monotonous, or chaotic regions of the model.

Using the data from one location in the ISD (Piazza San Marco) as the test distribution, the D_{BKS} statistic is calculated for each of the target distributions defined above, which is depicted in Table II.

E. Calculate the SPI score

The final step is to convert D_{BKS} into a more interpretable form to use as a comparison across soundscapes. Because the KS distance is a measure of dissimilarity, we first subtract it from one to give a measure of similarity between the test distribution and target distribution. This is then scaled to produce a score which ranges from 0 to 100, giving the final SPI formula

$$SPI = 100(1 - D_{BKS}\{MSN_{test}, MSN_{tgt}\}). \tag{4}$$

To show the usefulness of the test-target paradigm, we calculated the SPIs for each of the three target distributions for all of the locations included in the ISD, as shown in Table III. As each location is now assigned a SPI, this makes it possible to effectively produce three separate rankings of soundscape quality for these locations, depending on which target is considered the goal.

IV. EXPANDING THE SPI FRAMEWORK

Section III has defined and demonstrated the foundational methodology for calculating a SPI score. This

TABLE II. Two-dimensional K-S distance comparing the empirical test distribution (Piazza San Marco) against three soundscape target distributions.

Target	D_{BKS}
tgt ₁	0.64
tgt ₂	0.82
tgt ₃	0.28

TABLE III. SPI scores and rankings for the soundscapes of locations included in the ISD.

Ranking	SPI ₁ (pleasant)		SPI ₂ (calm)		SPI ₃ (vibrant)	
1	72	CarloV	64	CampoPrincipe	72	SanMarco
2	70	RegentsParkFields	54	CarloV	62	TateModern
3	64	CampoPrincipe	51	RegentsParkFields	60	StPaulsCross
4	64	RegentsParkJapan	50	PlazaBibRambla	59	Noorderplantsoen
5	63	PlazaBibRambla	46	MarchmontGarden	54	PancrasLock
6	62	RussellSq	44	MonumentoGaribaldi	54	TorringtonSq
7	62	MarchmontGarden	41	RussellSq	47	StPaulsRow
8	61	PancrasLock	40	PancrasLock	47	RussellSq
9	61	MonumentoGaribaldi	39	RegentsParkJapan	46	MiradorSanNicolas
10	55	StPaulsCross	32	StPaulsCross	42	CamdenTown
11	50	TateModern	32	MiradorSanNicolas	40	CarloV
12	50	StPaulsRow	31	TateModern	36	MonumentoGaribaldi
13	46	MiradorSanNicolas	30	TorringtonSq	34	MarchmontGarden
14	41	Noorderplantsoen	29	StPaulsRow	33	PlazaBibRambla
15	39	TorringtonSq	17	SanMarco	32	CampoPrincipe
16	35	SanMarco	16	Noorderplantsoen	31	EustonTap
17	22	CamdenTown	14	EustonTap	27	RegentsParkFields
18	17	EustonTap	14	CamdenTown	27	RegentsParkJapan

included how to define and sample a target distribution, prepare the test and target distributions for comparison using the KS distance metric, and convert this into a SPI score. To expand this methodology into an applicable framework, we define two distinct types of targets: bespoke targets and reference targets, each serving a unique purpose in the index development process.

A. Bespoke targets

Bespoke targets are essentially a direct application of the foundational method described above. Bespoke targets are tailor-made for specific projects, reflecting the desired soundscape perception for a particular application. These targets can be defined by stakeholders, designers, policy-makers, or decision-makers based on their unique requirements, objectives, and constraints. This flexibility allows the SPI for a specific project to be tailored to the desire of the stakeholders for how that specific soundscape should function. It can also provide a consistent and quantifiable baseline for scenarios like a soundscape design contest wherein a target is specified and provided to all participants in the contest, and the winning proposal is the design with the highest SPI score when assessed against that target. Stakeholders could use various methods to decide on a target, subject to the requirements of their project or use case. For example, it could be co-created with other stakeholders or space users, based on trying to match the soundscape of a previous project, or entirely arbitrary.

B. Reference targets

In contrast to bespoke targets, reference targets represent generalised, widely recognized soundscape archetypes which transcend specific applications or projects. These archetypes serve as reference points and enable comparisons across different domains and use cases. Essentially, a

reference target is a target that has been *empirically defined* to encapsulate the ideal of a particular type of soundscape (e.g., for a park, an urban square, a particular group of users, etc.).

1. Deriving a target based on a priori rankings

Absent from the above methodology has been an exploration of how to actually arrive at a target based on empirical evidence; i.e., not a target specified *ad hoc* but rather an “absolute” target, based on prior obtained evidence. Although arbitrary targets make the SPI framework incredibly flexible and able to score against an effectively infinite set of design goals, often targets should have some sort of systematic foundation, especially when defining a reference target.

The core challenge in developing a reference SPI target is determining what constitutes an ideal soundscape perception distribution for a given context. Whereas we can directly specify MSN parameters to create bespoke targets based on theoretical expectations or design goals, developing empirically grounded reference targets requires a more systematic approach.

Therefore, to enable this approach, we present one method of systematically deriving a target distribution based on a given ranking of soundscape quality. The *a priori* ranking serves as a bridge between existing knowledge about soundscape quality and the mathematical framework of the SPI. By starting with a ranking of soundscapes whose relative quality has been assessed through some external measure, we can use optimisation techniques to derive MSN parameters that

- (1) when used as a SPI target, produce scores that result in the same ranking order;
- (2) generate high SPI scores for the highly ranked soundscapes; and

(3) define a distribution in the circumplex space that captures the perceptual characteristics common to high-quality soundscapes in this context.

This approach allows us to work backward from known good (and poor) examples to define what the target distribution should look like. For instance, if we know that location *A* has a better soundscape than location *B* for our purposes, the optimal target distribution should result in location *A* receiving a higher SPI score than location *B*.

In this case study, we will examine a possible ranking from the ISD park locations produced by the authors (displayed in Table IV).

Effectively, this is an optimisation task to determine the MSN parameters which best achieve the above goals. Parameter optimisation refers to the process of adjusting the parameters of a system, model, or algorithm to achieve the best possible performance according to one or more objectives. To set up the optimisation task, we first need to express the parameter space and any constraints. Because our goal is to identify an optimised soundscape target distribution, the parameters that we will search over are

- $\xi = (\xi_x, \xi_y), -1 \leq \xi \leq 1;$
- $\Omega = \begin{pmatrix} \text{var}(x) & \text{cov}(x, y) \\ \text{cov}(y, x) & \text{var}(y) \end{pmatrix},$
 $- 0 \leq \text{var}(\cdot) \leq 1;$
 $- -1 \leq \text{cov}(\cdot) \leq 1;$
 $- \Omega$ must be symmetric and positive definite; and
- $\alpha = (\alpha_x, \alpha_y), -50 \leq \alpha \leq 50.$

Then, we define the objective functions based on the two goals given above. For each step in the algorithm with a given trial set of parameters, a target distribution will be produced, the SPI for each test location assessed according to the protocol described in Sec. III, and the resulting set of SPI scores and ranking will be scored using the objective functions. Goal (1) is assessed by calculating the Spearman rank correlation between the *a priori* ranking and SPI ranking such that

$$f_1 = r_s(R(\text{prior}), R(\text{target})). \tag{5}$$

TABLE IV. A predefined ranking of soundscape quality of the park locations included in the ISD. A SPI target will be derived which aims to reproduce this same ranking when applied to circumplex data from these locations.

Rank	Location
1	RegentsParkJapan
2	RegentsParkFields
3	CampoPrincipe
4	MonumentoGaribaldi
5	RussellSq
6	MiradorSanNicolas
7	StPaulsCross
8	Noorderplantsoen

Goal (2) is scored by calculating a weighted sum of the produced SPIs. To prioritize a target which provides high SPI scores for highly ranked soundscapes, we weight according to the ranking position,

$$f_2 = \sum_{i=1}^m \frac{1}{\text{rank}_i} \text{SPI}_i, \tag{6}$$

where *m* is the number of included locations, SPI_{*i*} is the calculated SPI score for the *i* th location assessed against a trial target, and rank_{*i*} is the calculated rank value of the *i* th location.

Through our testing, optimising only on the rank correlation regularly produced targets, which, although they did result in the desired ranking, were in no way representative of the soundscapes in question. We, therefore, aim to optimise for a consistent soundscape ranking and high SPI score for the top-ranked soundscapes. Optimising these parameters with respect to multiple objectives ensures a more holistic approach to system improvement, acknowledging the trade-offs and interactions between different goals.

We apply the nondominated sorting genetic algorithm (NSGA-II; Deb *et al.*, 2002) to optimise our target distribution parameters. NSGA-II is a popular and efficient multi-objective evolutionary algorithm that is well-suited for problems with multiple, potentially conflicting objectives.

The algorithm works as follows:²

- (1) Initialize a population of candidate solutions, each representing a set of target distribution parameters (ξ, Ω, α);
- (2) evaluate each candidate solution using the two objective functions defined above;
- (3) perform nondominated sorting to rank the solutions based on Pareto dominance;
- (4) calculate crowding distance for each solution to maintain diversity in the population;
- (5) select parent solutions using tournament selection based on nondomination rank and crowding distance;
- (6) create offspring solutions using crossover and mutation operators, ensuring that the constraints on the parameters are maintained;
- (7) combine parent and offspring populations and select the best solutions to form the next generation; and
- (8) repeat steps 2–7 for a specified number of generations or until a termination criterion is met.

The NSGA-II algorithm is implemented using the Python library `pymoo v0.6.1.3` (Blank and Deb, 2020). The population size is set to 150, and the algorithm runs for 100 generations. In `pymoo`, each objective function is supposed to be minimised such that when implementing the algorithm and in the results, $-f_1$ and $-f_2$ are used. After running the NSGA-II algorithm, we obtain a set of non-dominated solutions representing the Pareto front, displayed in Fig. 3(a). Each solution on the Pareto front represents a trade-off between the two objectives: maximising the rank correlation (*f*₁) and maximising the weighted sum of SPI




FIG. 3. (Color online) NSGA-II optimisation to learn the MSN parameters which produce the Park ranking, which include (a) multi-objective optimisation Pareto front, where the selected solution is indicated in red, and (b) SCM distribution of the derived target distribution.

scores (f_2). The Pareto front allows us to visualize and analyse the range of possible solutions from those that prioritize ranking consistency to those that emphasise high-SPI scores for top-ranked soundscapes.

The Pareto solution which minimises the Euclidean distance from the ideal point ($\min -f_1, \min -f_2$) is extracted. This approach provides a balance between ranking consistency and high SPI scores for top-ranked soundscapes. Once the optimal solution is selected, we can sample from the MSN distribution and plot the derived target distribution, which is shown in Fig. 3(b):

$$\text{tgt}_{\text{park}} \sim \begin{cases} \xi = [0.694, 0.406], \\ \Omega = \begin{bmatrix} 0.157 & 0.040 \\ 0.040 & 0.255 \end{bmatrix}, \\ \alpha = [5.054, -37.671]. \end{cases} \quad (7)$$

The resulting tgt_{park} , with the MSN parameters given above, exhibits some expected characteristics: it is almost entirely pleasant with a long uneventful tail into the calm quadrant, but, somewhat unexpectedly, the mode is slightly vibrant. What should be made clear about this demonstration is that we are not presenting this as a reference target to be used in the future—this section is meant as a demonstration of a method which can be used to derive a target from an *a priori* ranking. In this case, the *a priori* ranking was created by the authors from their experience of the locations in the ISD. To truly be called an empirical reference target, the ranking would need to be arrived at empirically via some other metric (e.g., health or productivity ratings of the areas) or through an experiment such as paired-choice comparisons.

It is these reference targets with an empirical backing which would ideally form agreed on standards and benchmarks in the field against which new soundscapes would be compared. The best methods for empirically determining the ideal soundscape distribution for a given context will no doubt remain a topic of debate and development in the coming years.

V. DISCUSSION

The development of bespoke and reference context-dependent SPIs represents a significant step toward enabling more comprehensive and effective applications of the soundscape approach. By providing a unified framework for defining these indices, the potential for quantifying and comparing soundscape quality across diverse contexts and applications is unlocked while still ensuring that the multidimensional and context-driven aspects of soundscape quality are considered.

A. Applications of the SPI framework

The proposed framework offers several key advantages. First, it acknowledges the inherent context-dependent nature of soundscape perception, allowing for the creation of indices tailored to specific use cases or design goals through the use of bespoke targets. This flexibility ensures that the resulting SPIs accurately capture the desired soundscape perception for the given application, enabling targeted interventions and optimisations.

Second, the inclusion of reference targets facilitates cross-comparisons and benchmarking, enabling a common language and understanding of soundscape quality across different domains. By calculating the distance between a given soundscape and these widely recognized references, stakeholders can identify areas for improvement and prioritize interventions accordingly, aligning their efforts with collectively recognized standards of desirable or undesirable soundscapes.

We expect that this would then expand into collections of SPI targets. As an example, imagine trying to define a SPI that could be applied across an entire city. A single index is insufficient because each type of place within the city (e.g., parks, plazas, and residential areas) has different requirements for its soundscape. Therefore, each place type would need its own soundscape target.

In this example, these sets of targets would correspond to different types of places within the city (e.g., a single

target for parks, a target for plazas, etc.). When applying this “urban typology” set of targets, the soundscape of each location being assessed would be scored against its relevant target (i.e., how well does a specific park perform in comparison to a reference park target). This results in a single score for each location that can be compared against all other locations, regardless of whether or not they are the same type of place, allowing for different soundscapes to be compared on a common scale. This system ensures that context (in this case, the typology of a space) is brought into the assessment, permitting soundscapes to be scored against the most appropriate target. Enabling these context-dependent assessments to be expressed on a common scale can facilitate additional use cases such as soundscape mapping, which requires a single scale to be applied across an entire city.

This set of targets, made up of different space typologies, is just one example of an application of reference SPIs. Other examples could include a demographics SPI, where different targets are set for respondents from different demographic groups or a “use case” SPI with different targets set for different intended purposes of spaces (e.g., recreation, restoration, and socialising). We encourage users of the SPI to define their own single reference targets that can be added to these suites of targets for use by others and their own new sets of references.

Kogan *et al.* (2018, Fig. 6), in fact, displays a startlingly similar concept, showing the locations of the three categories of traffic noise dominance (“traffic noise,” “balanced,” and “natural”) plotted in the circumplex perceptual model. It can be clearly observed in this plot that the GSI categories create their own clusters within the circumplex.

Although it is expected that the target distribution would usually represent the ideal or goal soundscape perception, it is also possible to define target distributions that represent undesirable or suboptimal soundscape perceptions. For instance, in a soundscape mapping context, it may be beneficial to map and identify chaotic soundscapes across a city to better target areas for soundscape interventions. In this case, the target distribution would be set in the chaotic quadrant, and a higher SPI would indicate a closer alignment with the target distribution. This flexibility allows the SPI to be applied to a wide range of contexts and applications, enabling the quantification and comparison of soundscape quality across diverse scenarios.

B. Connecting with soundscape indicators and acoustic data

Unlike previous soundscape indices (see Sec. II A), SPI does not include any direct connection to soundscape indicators such as the sound level, spectral content, etc. Its basis in perceptual descriptor data effectively allows the analysis and quantification of soundscape information to be modularised, separating the task of calculating a single index from the complex task of predicting soundscape perception from objective data. The modularisation of soundscape data analysis allows the entire pipeline from environmental data collection through to soundscape index scoring to remain

flexible. Following the soundscape engineering paradigm exhibited in Mitchell *et al.* (2023), the connection between soundscape indicators through descriptors to indices can be made by predictive soundscape models. These models are trained on increasingly large scale datasets and generally designed to predict soundscape descriptors, including the SCM attributes (Hou *et al.*, 2024; Ooi *et al.*, 2022). With the complex and multidimensional nature of soundscape perception and rapid progression in machine learning techniques and applications, an index framework should be able to integrate new and improved models. By separating the prediction of perceptual descriptors based on objective metrics from the calculation of the single value index itself, the SPI framework allows for these predictive models and the creation of new indices to advance independently.

C. Comparison with existing soundscape indices

The SPI framework represents a unique approach to soundscape assessment, building on and differentiating itself from previous indices in several key ways. First, the SPI framework is fundamentally perception focussed. By referring to the soundscape perception index, we aim to highlight the unique and perception-focussed nature of this index framework. The SPI core method operates entirely within the perception data space with no direct reference to acoustic or other indicators. Perceptual data (or predicted perceptual data) are the only operant factors of the SPI method. The aim of SPI is to combine multidimensional perception data and context (including design goals) into a single metric.

The SPI framework shares this important characteristic with indices like the SDI (Liu *et al.*, 2014) and SHD (Guo *et al.*, 2023) in that they all prioritize perception as the primary input for assessment. SDI and SHD, however, focus on the perception of sound sources that can be observed in an acoustic environment and their relationships. Although these indices offer valuable insights into specific aspects of soundscape perception, they are somewhat limited in their scope and adaptability. The SPI framework builds on these efforts by incorporating the full dimensionality of the SCM and allowing for context-sensitive assessments through bespoke and reference targets. This approach enables the SPI to address a wider range of soundscape evaluation needs while maintaining the crucial focus on perceptual data that distinguishes these methods from purely acoustic measurements.

Furthermore, the SPI framework is designed to be generalisable, extensible, and adaptable. Unlike previous indices that often represent a single target in a particular context, the SPI framework allows for scoring soundscapes against any goal defined by the user. This flexibility makes it applicable across a wide range of contexts and design objectives, from urban planning and acoustic design to research and policy development.

D. Considerations and future work

Several considerations should be noted when defining a SPI target. First, the target distribution should be

representative of the desired soundscape perception for the given application. This requires a clear understanding of the context, objectives, and constraints of the project, as well as the preferences and expectations of stakeholders and end users. Second, the temporal and spatial scales of the target distribution should align with the soundscape assessment being conducted. What constitutes the actual spatial bounds of “a soundscape” or, indeed, “a place” is a complex question which will depend on the context of the assessment. For example, a park soundscape may be defined by the boundaries of the park itself, or it may extend to include the surrounding urban environment or be restricted to a certain distance from a feature of interest in the park.³ The temporal scale of the assessment is also important as soundscape quality can vary throughout the day and across different seasons. Increasing the spatial bounds of what is considered to be the soundscape under examination (e.g., a single position vs a 25 m² area vs an entire park) or extending the temporal scale will almost certainly result in a distribution with a larger variance. What scales are appropriate for a given assessment will depend on the context and objectives of the project, but they should be considered when defining the target distribution. Applying a target distribution that is too broad or too narrow for the context of the assessment may result in inaccurate or misleading SPI scores. As the circumplex distribution first described in [Mitchell et al. \(2022\)](#) and further formalised here develops, we are hopeful that a better understanding of the relationship between temporal and spatial scales and the parameters of the distribution will emerge and contribute to an increased understanding of what constitutes “the soundscape of a place” and how this should be reflected in its ideal perception distribution.

Various other distance metrics were considered when developing the SPI method. The simplest method is to define a single point target rather than a target distribution and calculate a normalised mean Euclidean distance between points in the test distribution and target point. Although this is conceptually simple and requires defining only a single coordinate point as a target rather than the MSN parameters described in [Sec. III B](#), the shape and spread of a soundscape distribution is itself an important factor in describing the collective perception of a soundscape and would not be captured by this method.

An additional method that was considered was to regard a target as an ellipse (or, indeed, any other shape) drawn in the circumplex space (similar to the simplified median decile curves proposed in [Mitchell et al., 2022](#)). A SPI score would then be calculated based on the percentage of responses which fall within the space defined by the ellipse. Again, this is conceptually quite simple, and defining the ellipse targets is straightforward. However, this method has an important flaw—it is easy to artificially inflate or deflate the scores merely by changing the area of the ellipse. The larger the ellipse, the higher all SPI scores will be regardless of whether the sample distribution is wide or narrow. This would also limit cross-comparability between targets. As can be noticed in [Table III](#), defining a target distribution

with a larger spread (i.e., tgt_1) does not automatically result in higher SPI scores across the board as it would with the ellipse target method. By defining the SPI as a true target-test distribution comparison, we ensure that the SPI always accurately reflects the similarity between the perception of a soundscape and its target, in terms of its location in the circumplex and the shape of the data.

As noted in [Sec. IV B 1](#), although a methodology for deriving targets is presented, the *a priori* ranking that we use for the demonstration was not itself arrived at empirically. Hence, the park target cannot be considered to be a true reference target. A key piece of future work is to use experimental methods, such as paired-choice comparisons, to arrive at a well-defined ranking, which can then produce a true reference target.

VI. CONCLUSION

The European Research Council (ERC)-funded Soundscape Indices (SSID) Project was started mostly with the ambition to derive soundscape indices that could serve as numerical/quantifiable tools ([Kang et al., 2019](#)) to better inform urban sound planning and design decisions. Any soundscape researcher having ever made an attempt at defining “the” soundscape quality index will know what a challenging, even impossible task this is. Some may even argue that trying to reduce soundscape quality to a single value quantity, and deriving any soundscape index, could be what philosophers would call a *contradictio in adjecto* as the soundscape approach intrinsically advocates for a multidimensional characterisation of the acoustic environments that we experience in our lives. For these reasons, we felt that it was necessary to take a step back and create, instead, a framework tailored for the field specifically that could easily be adapted to different contexts and capture the multifaceted aspects of the soundscape of a place.

The proposed framework addresses the existing gap in quantifying multidimensional soundscape perception, facilitating a broader application of the soundscape approach in areas such as urban planning, environmental management, acoustic design, and policy development. Through the creation of bespoke indices tailored to specific design goals and use of reference targets for benchmarking, this framework empowers stakeholders and decision-makers to make informed choices and prioritize soundscape improvements aligned with their unique objectives and constraints.

Furthermore, the grounding of the SPI framework in the SCM ensures a robust theoretical foundation, capturing the multidimensional nature of soundscape perception. The use of a distance metric enables quantitative assessments and comparisons, fostering a common language and understanding of soundscape quality across different domains. This shared understanding facilitates knowledge exchange, collaborative efforts, and the development of best practices within the field. As the SPI framework continues to be explored and refined, future research should focus on validating and expanding the range of reference targets as well

as investigating the potential for incorporating additional dimensions and factors that influence soundscape perception. The integration of emerging technologies (such as virtual, mixed, and augmented reality) may also provide new avenues for immersive soundscape evaluation and index development. Additionally, the application of the framework in diverse real-world scenarios, ranging from urban planning and environmental management to acoustic design and policy development, will provide valuable insight and contribute to the ongoing refinement and adaptation of the SPI framework.

In many ways, the proposed SPI framework is not so conceptually different from the whole idea of dB-based set of indicators that the SSIDs project itself is trying to “overcome.” There is no such thing as a single noise indicator (L) to univocally describe sound levels in all circumstances; rather, different noise indicators are defined for different scenarios and temporal or spectral requirements (e.g., L_{den} , $L_{Aeq,T}$, etc.), based on testing needs. The dB is the unit for all of them, but A -weighted equivalent sound levels for a 1-h interval cannot be directly compared with whole-day indicators with penalties. We are trying to achieve the same with the SPI to provide a way of defining different indices for different contexts while maintaining a consistent framework.

Ultimately, for the SPI approach to succeed, collaboration with stakeholders, end users, and experts from various domains will be crucial in ensuring the framework’s relevance and applicability across a wide range of contexts.

SUPPLEMENTARY MATERIAL

See the supplementary material for further technical details of the multi-objective optimisation procedure and its associated code.

ACKNOWLEDGMENTS

We would like to thank Dr. Giorgios Petrou for his expert advice and guidance on multi-objective optimisation methods. We would also like to thank Dr. Nicole Watson for her advice on the structure and language of the paper. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant No. 740696, project title Soundscape Indices—SSID). More information and related publications can be found at the CORDIS webpage of the project.⁴ The authors contributed accordingly: A.M., conceptualization, methodology, software, writing—original draft, writing—review and editing, and visualization; F.A., conceptualization, methodology, and writing—review and editing; T.O., conceptualization and writing—review and editing; and J.K., conceptualization, supervision, and funding acquisition.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are openly available in ISD (version 1.0.1-alpha.1; Mitchell *et al.*, 2024) on Zenodo at <https://doi.org/10.5281/zenodo.10672568>. The code to recreate the analysis in this paper can be found on this paper’s OSF page (Mitchell, 2024) at <https://doi.org/10.17605/OSF.IO/Q3M5X>.

¹It is important to note that the parameters which appear in the density expression (ξ, Ω, α) are called direct parameters (DP). They directly parameterise a MSN density and are typically only estimated by fitting a MSN to a sample. The more familiar and interpretable components (mean, standard deviation, and skewness) are termed the centred parameters (CP). It is possible to move from one parameterisation to another, however “while any choice of the DP components is admissible, the same is not true for CP”; i.e., we can always move DP \rightarrow CP but not always CP \rightarrow DP. In this context, it is most important for readers not to confuse the location parameter ξ with the sample mean μ . A more complete explanation of these parameterisations can be found in Azzalini (2016).

²Further technical details of the multi-objective optimisation procedure and relevant code can be found in the supplementary material.

³For instance, the SSID protocol, which produced the data used in this paper, attempted to address this by considering its spatial bounds for what constitutes one location to be “an ‘environmental unit,’ wherein the environmental factors are consistent, and is typically perceived to constitute a single distinct area,” noting that “the exact dimensions and delineation of the environmental unit will vary depending on the characteristics of the space” (Mitchell *et al.*, 2020).

⁴See <https://cordis.europa.eu/project/id/740696/factsheet> (Last viewed 2024-5-28).

Aletta, F., and Kang, J. (2015). “Soundscape approach integrating noise mapping techniques: A case study in Brighton, UK,” *Noise Mapp.* 2(1), 1–12.

Aletta, F., Kang, J., and Axelsson, Ö. (2016). “Soundscape descriptors and a conceptual framework for developing predictive soundscape models,” *Landscape Urban Plann.* 149, 65–74.

Aletta, F., Mitchell, A., Oberman, T., Kang, J., Khelil, S., Bouzir, T. A. K., Berkouk, D., Xie, H., Zhang, Y., Zhang, R., Yang, X., Li, M., Jambrošić, K., Xaninović, T., van den Bosch, K., Lühr, T., Orlik, N., Fitzpatrick, D., Sarampalis, A., Aumond, P., Lavandier, C., Moshona, C. C., Lepa, S., Fiebig, A., Papadakis, N. M., Stavroulakis, G. E., Sudarsono, A. S., Sarwono, S. J., Puglisi, G. E., Jafari, F., Astolfi, A., Shtrepi, L., Nagahata, K., Jo, H. I., Jeon, J. Y., Lam, B., Chieng, J., Ooi, K., Hong, J. Y., Antunes, S. M., Alves, S., de Ulhoa Carvalho, M. L., Michalski, R. L. X. N., Kogan, P., Manzano, J. V., Quesada, R. G., Silva, E. S., Pastor, J. A. A., Nilsson, M. E., Axelsson, Ö., Gan, W.-S., Watcharasapat, K. N., Jaratjarungkiat, S., Ong, Z.-T., Yorukoglu, P. N. D., Osma, U. B. E., and Nguyen, T. L. (2024). “Soundscape descriptors in eighteen languages: Translation and validation through listening experiments,” *Appl. Acoust.* 224, 110109.

Aletta, F., Oberman, T., and Kang, J. (2018). “Associations between positive health-related effects and soundscapes perceptual constructs: A systematic review,” *Int. J. Environ. Res. Public Health* 15(11), 2392.

Aletta, F., and Torresin, S. (2023). “Adoption of ISO/TS 12913-2:2018 protocols for data collection from individuals in soundscape studies: An overview of the literature,” *Curr. Pollut. Rep.* 9, 710–723.

Axelsson, Ö. (2015). “How to measure soundscape quality,” in *Proceedings of Euronoise 2015*, Maastricht, Netherlands June 1–3 [Nederlands Akoestisch Genootschap and ABAV (Belgian Acoustical Society)], pp. 1477–1481.

Axelsson, Ö., Nilsson, M. E., and Berglund, B. (2010). “A principal components model of soundscape perception,” *J. Acoust. Soc. Am.* 128(5), 2836–2846.

Axelsson, Ö., Nilsson, M. E., and Berglund, B. (2012). “The Swedish soundscape-quality protocol,” *J. Acoust. Soc. Am.* 131, 3476.

Azzalini, A. (2005). “The skew-normal distribution and related multivariate families,” *Scand. J. Statist.* 32(2), 159–188.

- Azzalini, A. (2016). "How to sample from the SN and related distributions when we want to fix skewness and other cumulants," available at http://azzalini.stat.unipd.it/SN/how_to_sample.pdf (Last viewed November 17, 2024).
- Azzalini, A. (2021). "The R package sn: The skew-normal and related distributions such as the skew- t and the SUN," available at <https://cran.r-project.org/package=sn> (Last viewed November 17, 2024).
- Azzalini, A., and Capitanio, A. (1999). "Statistical applications of the multivariate skew normal distribution," *J. R. Statist. Soc. Ser. B* **61**(3), 579–602.
- Azzalini, A., and Valle, A. D. (1996). "The multivariate skew-normal distribution," *Biometrika* **83**(4), 715–726.
- Berglund, B., Lindvall, T., and Schwela, D. H. (1999). *Guidelines for Community Noise* (World Health Organization, Geneva).
- Blank, J., and Deb, K. (2020). "Pymoo: Multi-objective optimization in python," *IEEE Access* **8**, 89497–89509.
- Blauert, J., and Jekosch, U. (1997). "Sound-quality evaluation—A multi-layered problem," *Acta Acust. Acust.* **83**(5), 747–753.
- Cao, X., Meng, Q., and Kang, J. (2020). "Red soundscape index (RSI): An index with the potential to assess soundscape quality," in *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, Seoul, August 23–26 (INCE–USA, Wakefield, MA), Vol. 261, pp. 3527–3539.
- Chakravati, I. M., Laha R. G., and Roy J. (1967). *Handbook of Methods of Applied Statistics* (Wiley, New York), Vol. 1, pp. 392–394.
- Chen, X., and Kang, J. (2023). "Natural sounds can encourage social interactions in urban parks," *Landscape Urban Plann.* **239**, 104870.
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.* **6**(2), 182–197.
- European Union (2002). *European Parliament and Council, Directive 2002/49/EC Relating to the Assessment and Management of Environmental Noise* (Publications Office of the European Union, Brussels).
- Fasano, G., and Franceschini, A. (1987). "A multidimensional version of the Kolmogorov–Smirnov test," *Mon. Not. R. Astron. Soc.* **225**(1), 155–170.
- Fastl, H. (2006). "Psychoacoustic basis of sound quality evaluation and sound engineering," in *The Thirteenth International Congress on Sound and Vibration*, Vienna, Austria, July 2–26 (IIAV, Auburn, AL).
- Fiebig, A. (2018). "Does it make a difference to have soundscape standards?," *Proceedings—Euronoise 2018*, Crete, May 27–31 (EEA, Helina), 6 pp., available at https://www.euronoise2018.eu/docs/papers/482_Euronoise2018.pdf (Last viewed November 17, 2024).
- Fletcher, H., and Munson, W. A. (1933). "Loudness, its definition, measurement and calculation," *Bell Syst. Tech. J.* **12**(4), 377–430.
- Guo, X., Liu, J., Chen, Z., and Hong, X.-C. (2023). "Harmonious degree of sound sources influencing visiting experience in Kulangsu Scenic Area, China," *Forests* **14**(1), 138.
- Guski, R. (1997). "Psychological methods for evaluating sound quality and assessing acoustic information," *Acta Acust. Acust.* **83**(5), 765–774.
- Hellman, R., and Zwicker, E. (1987). "Why can a decrease in dB(A) produce an increase in loudness?," *J. Acoust. Soc. Am.* **82**(5), 1700–1705.
- Hou, Y., Ren, Q., Mitchell, A., Wang, W., Kang, J., Belpaeme, T., and Botteldooren, D. (2024). "Soundscape captioning using sound affective quality network and large language model," [arXiv:2406.05914](https://arxiv.org/abs/2406.05914).
- ISO/TS 12913-2 (2018). "Acoustics—Soundscape—Part 2: Data collection and reporting requirements" (International Organization for Standardization, Geneva, Switzerland).
- ISO/TS 12913-3 (2019). "Acoustics—Soundscape—Part 3: Data analysis" (International Organization for Standardization, Geneva, Switzerland).
- Kang, J. (2006). *Urban Sound Environment* (Taylor & Francis, London, UK).
- Kang, J. (2023). "Soundscape in city and built environment: Current developments and design potentials," *City Built Environ.* **1**(1), 1.
- Kang, J., and Aletta, F. (2018). "The impact and outreach of soundscape research," *Environments* **5**(5), 58.
- Kang, J., Aletta, F., Oberman, T., Erfanian, M., Kachlicka, M., Lionello, M., and Mitchell, A. (2019). "Towards soundscape indices," in *Proceedings of the 23rd International Congress on Acoustics*, 9–13 September, Aachen, Germany (EAA Euregio, Aachen, Germany), pp. 2488–2495.
- Kang, J., Aletta, F., Oberman, T., Mitchell, A., Erfanian, M., Tong, H., Torresin, S., Xu, C., Yang, T., and Chen, X. (2023). "Supportive soundscapes are crucial for sustainable environments," *Sci. Total Environ.* **855**, 158868.
- Kang, J., and Schulte-Fortkamp, B. (2016). *Soundscape and the Built Environment* (CRC Press, Boca Raton, FL), pp. 197–214.
- Kogan, P., Arenas, J. P., Bermejo, F., Hinalaf, M., and Turra, B. (2018). "A green soundscape index (GSI): The potential of assessing the perceived balance between natural sound and traffic noise," *Sci. Total Environ.* **642**, 463–472.
- Kryter, K. D. (1994). *The Handbook of Hearing and the Effects of Noise* (Academic, London, UK).
- Lionello, M., Aletta, F., and Kang, J. (2020). "A systematic review of prediction models for the experience of urban soundscapes," *Appl. Acoust.* **170**, 107479.
- Liu, J., Kang, J., Behm, H., and Luo, T. (2014). "Effects of landscape on soundscape perception: Soundwalks in city parks," *Landscape Urban Plann.* **123**, 30–40.
- Mitchell, A. (2022). "Predictive modelling of complex urban soundscapes: Enabling an engineering approach to soundscape design," Ph.D. thesis, University College London, available at <https://discovery.ucl.ac.uk/id/eprint/10156562> (Last viewed November 17, 2024).
- Mitchell, A. (2024). Soundscape Perception Indices (SPI). <https://doi.org/10.17605/OSF.IO/Q3M5X>
- Mitchell, A., Aletta, F., and Kang, J. (2022). "How to analyse and represent quantitative soundscape data," *JASA Express Lett.* **2**(3), 037201.
- Mitchell, A., Aletta, F., Oberman, T., Erfanian, M., and Kang, J. (2023). "A conceptual framework for the practical use of predictive models and Soundscape Indices: Goals, constraints, and applications," in *INTER-NOISE 2023 Conference*, August 20–23, Chiba, Greater Tokyo (INCE–USA, Wakefield, MA).
- Mitchell, A., Oberman, T., Aletta, F., Erfanian, M., Kachlicka, M., Lionello, M., Fang, X., and Kang, J. (2024). "The International Soundscape Database: An integrated multimedia database of urban soundscape surveys – questionnaires with acoustical and contextual information," available at <https://doi.org/10.5281/zenodo.10672568> (Last viewed November 17, 2024).
- Mitchell, A., Oberman, T., Aletta, F., Erfanian, M., Kachlicka, M., Lionello, M., and Kang, J. (2020). "The Soundscape Indices (SSID) Protocol: A method for urban soundscape surveys—Questionnaires with acoustical and contextual information," *Appl. Sci.* **10**(7), 2397.
- Ooi, K., Watcharasupat, K. N., Lam, B., Ong, Z.-T., and Gan, W.-S. (2022). "Probably pleasant? A neural-probabilistic approach to automatic masker selection for urban soundscape augmentation," in *ICASSP 2022—2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 23–27 May, Singapore (IEEE, New York).
- Parmanen, J. (2007). "A-weighted sound pressure level as a loudness/annoyance indicator for environmental sounds—Could it be improved?," *Appl. Acoust.* **68**(1), 58–70.
- R Core Team. (2018). *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna, Austria), available at <https://www.R-project.org/> (Last viewed November 17, 2024).
- Russell, J. A. (1980). "A circumplex model of affect," *J. Pers. Soc. Psychol.* **39**(6), 1161–1178.
- Schulte-Fortkamp, B., Fiebig, A., Sisneros, J. A., Popper, A. N., and Fay, R. R., eds. (2023). *Soundscapes: Humans and Their Acoustic Environment* (Springer International, Cham, Switzerland).
- Schulte-Fortkamp, B., and Kang, J. (2013). "Introduction to the special issue on soundscapes," *J. Acoust. Soc. Am.* **134**(1), 765–766.
- Xiang, Y., Meng, Q., Zhang, X., Li, M., Yang, D., and Wu, Y. (2023). "Soundscape diversity: Evaluation indices of the sound environment in urban green spaces—Effectiveness, role, and interpretation," *Ecol. Indic.* **154**, 110725.
- Xu, Z.-y., Chen, L., Pijanowski, B. C., and Zhao, Z. (2023). "A frequency-dependent acoustic diversity index: A revision to a classic acoustic index for soundscape ecological research," *Ecol. Indic.* **155**, 110940.
- Yang, D., Cao, X., and Meng, Q. (2022). "Effects of a human sound-based index on the soundscapes of urban open spaces," *Sci. Total Environ.* **802**, 149869.
- Zwicker, E., and Fastl, H. (2007). *Psychoacoustics: Facts and Models*, 3rd ed. (Springer, Berlin), p. 462.