Repetition detection and rapid auditory learning for stochastic tone clouds

Running title: Rapid learning of tone clouds

Trevor R. Agus

SARC, School of Arts, English and Languages, Queen's University Belfast, Belfast, 1
Malone Road, BT7 1NN, United Kingdom.

Daniel Pressnitzer

Laboratoire des Systèmes Perceptifs, Département d'études cognitives, École Normale
Supérieure, PSL University, CNRS, 29 rue d'Ulm, 75005, Paris, France

Date: Wednesday, 7th June, 2021

1 ABSTRACT

2 Stochastic sounds are useful to probe auditory memory, as they require listeners to learn
3 unpredictable and novel patterns under controlled experimental conditions. Previous studies
4 using white noise or random click trains have demonstrated rapid auditory learning. Here, we
5 explored perceptual learning with a more parametrically variable stimulus. These “tone clouds”
6 were defined as broadband combinations of tone pips at randomized frequencies and onset
7 times. Varying the number of tones covered a perceptual range from individually audible pips
8 to noise-like stimuli. Results showed that listeners could detect and learn repeating patterns in
9 tone clouds. Task difficulty varied depending on the density of tone pips, with sparse tone
10 clouds the easiest. Rapid learning of individual tone clouds was observed for all densities, with
11 a roughly constant benefit of learning irrespective of baseline performance. Variations in task
12 difficulty were correlated to amplitude modulations in an auditory model. Tone clouds thus
13 provide a tool to probe auditory learning in a variety of task-difficulty settings, which could be
14 useful for clinical or neurophysiological studies. They also show that rapid auditory learning
15 operates over a wide range of spectrotemporal complexity, essentially from melodies to noise.

I. INTRODUCTION

One of the important roles of hearing is to recognize sounds that have been heard before, using memory to inform behavior. Recent psychophysical studies have demonstrated a rapid and powerful auditory learning mechanism, whereby listeners exhibit large changes in their performance patterns after only a few exposures to novel sounds. These studies used stochastic sounds: noise (Agus *et al.*, 2010; Agus and Pressnitzer, 2013; Luo *et al.*, 2013; Andrillon *et al.*, 2015; Viswanathan *et al.*, 2016; Song and Luo, 2017), random click trains (Kang *et al.*, 2017; Kang *et al.*, 2018) and discrete sequences of tones (Bianco *et al.*, 2020). Such sounds are useful in so far as they plausibly push any learning mechanism to the limit. Gaussian noise, for instance, is fully unpredictable, so it has maximum entropy. However, noise, click trains and melodic bursts are not necessarily representative of the wide range of spectrotemporal complexity¹ that is found in natural listening conditions (e.g., McDermott and Simoncelli, 2011).

The main purpose of the current study was to probe the same learning mechanism with another class of stochastic sounds, one which would allow for a fine parametric control of spectrotemporal complexity. There are convincing examples of perceptual learning for discrimination tasks of simple sounds (see Wright and Zhang, 2009, for a review), over long time scales for training (e.g., Lim and Holt, 2011), or in a supervised manner (e.g., Mishra and Panda, 2016), but, to the best of our knowledge, the rapid and unsupervised learning of sounds with an acoustic complexity comparable to natural sounds has not yet been demonstrated systematically. If rapid learning could be observed over a broad range of complexities, then this would support the notion that rapid learning is useful for learning about everyday sounds. Another purpose was methodological: if perceptual learning can be observed in a task whose

1 difficulty can be varied parametrically, the stimuli could be adapted to future
2 neurophysiological or clinical studies.

3 The behavioral paradigm used here follows Agus *et al.* (2010) and subsequent “memory for
4 noise” studies (e.g., Luo *et al.*, 2013; Viswanathan *et al.*, 2016). In this paradigm, listeners
5 were asked to detect seamless repetitions in Gaussian noises within many experimental trials.
6 Memory was probed by presenting the exact same snippet of noise on several trials, randomly
7 interspersed within an experimental block and without telling the listeners of the presence of
8 such trials. The hypothesis was that any learning of a re-occurring noise snippet *across* trials
9 would make the *within* trial repetition detection task easier. Learning was indeed observed, and
10 it was surprisingly rapid: when learning occurred, it only took a few exposures to build up
11 (Agus *et al.*, 2010). Learning was also observed when attention was diverted from the sounds
12 to be learnt (Andrillon *et al.*, 2015) or during REM sleep (Andrillon *et al.*, 2017). Whatever
13 the underlying mechanism for the memory of noise in such experiments, behavioral findings
14 show that it is capable of learning fully unpredictable sounds, rapidly, amongst statistically
15 identical samples, and it can operate without the need of focused attention. Such a mechanism
16 would thus possess many desirable features for the learning of everyday sounds.

17 One obvious disadvantage of using Gaussian noise, however, is that the within-trial repetition
18 detection task itself is relatively difficult, at least over the range of noise durations that have
19 been tested so far (0.125–4 s; Agus *et al.*, 2010). At the duration tested in the present study
20 (0.5 s), baseline sensitivity for repetition detection before learning averaged $d' = 0.5$ for naïve
21 listeners. This might make the task unsuitable for animal models or special populations with
22 restricted time for training. Furthermore, pilot data showed that detecting repetitions in noise
23 was difficult in noisy environments, such as an MRI scanner (Kumar *et al.*, 2014). Kang *et al.*
24 (2017) modulated the difficulty of the baseline task by controlling the average inter-click rate

1 within click trains. Rapid learning was observed for all baseline task difficulties, suggesting
2 that the paradigm remains relevant irrespective of baseline difficulty. However, random click
3 trains lack the spectrotemporal complexity of noise or most natural sounds.

4 Here, we explored a new class of stochastic stimuli to address those issues, a subset of which
5 we originally created to be presented with simultaneous recording in an MRI scanner (Kumar
6 *et al.*, 2014). We wished to parametrically control the complexity of the sounds, to make the
7 baseline repetition detection task more or less difficult as required, and also to probe a range
8 of spectrotemporal complexities more representative of the types of sound that may be relevant
9 to a listener outside of the laboratory. Also, to ensure that repetitions were detected solely from
10 the repetitive nature of the stimuli, it was necessary to generate repeated and unrepeated
11 versions of the stimuli without trivial local cues at the repetition boundary, such as a click or a
12 statistical discontinuity. With these requirements, “tone clouds” were designed as the sum of
13 many short pure-tone pips, distributed in a pseudorandom fashion over time and frequency.
14 The distribution was specified in such a way that the local statistical properties of longer
15 segments (e.g., those used to generate unrepeated clouds) would be identical to those of shorter
16 segments (used to generate repeated clouds). This was achieved by imposing a “grid” in time
17 and log-frequency, such that one and only one tone pip was guaranteed to start in each time/log-
18 frequency cell of the grid (Figure 1). The tone pips that started late enough in their cell were
19 allowed to tail over into the next cell. This applied equally to the grid line that marked the end
20 of one segment and the start of its repetition so that there was no silence nor statistical
21 discontinuity between repeated segments, and thus no local cues to within-trial repetition.

Figure 1. Schematics and cochleagrams of some example tone clouds. **A:** Each black rectangle represents a tone pip, with gray level indicating amplitude (with particularly light gray at the limits of frequency). The red lines indicate the time-frequency grid with which these tone clouds were designed, here with parameters of 2 columns per second and 1 row per octave. Exactly one tone pip starts in each cell of the grid, but because of the random offset position within its cell, its offset may stray into the subsequent cell. **B:** Cochleagram of the stimulus outlined in Panel A, with local intensity represented as brightness. **C and D:** As for panels A and B, but for a denser stimulus generated with more columns/s. **E and F:** As for panels C and D, but for a yet denser stimulus generated with more rows/octave. **G:** Cochleagram of a noise spectrally matched to the tone clouds. **H:** Cochleagram of one of the densest tone clouds used in the experiments (512 columns/s × 256 rows/octave).

The tone clouds as defined here thus retain features of Gaussian noise that make them suitable for learning tasks using repetition detection as a baseline: the stimuli are stochastic, broadband, complex, and listeners will not have heard them before the start of the experiment. A novelty is that we could now probe learning for a broad range of spectrotemporal complexities, spanning a perceptual space ranging from random “melodies”, for sparse tone clouds with relatively few tone pips (cf. Bianco *et al.*, 2020), to denser tone clouds with more tone pips, reminiscent of informational masking studies (Gutschalk *et al.*, 2008), and up to the densest tone clouds that were essentially indistinguishable from stationary noise (cf. Rice, 1952).

In the preliminary Experiment 1, a task similar to the classic repetition-detection paradigm introduced by Guttman and Julesz (1963) was used to measure listeners’ baseline ability to detect repetitions, with the important difference that we used tone clouds of different densities instead of Gaussian noise. Then, the main Experiment 2 probed whether perceptual learning was observed for tone clouds across a wide range of densities. The task in this second experiment was identical to Agus *et al.* (2010), here using tone clouds again instead of noise (Agus *et al.*, 2010) or random click trains (Kang *et al.*, 2017). Finally, a simple functional model based on the presence or absence of local energy variations after cochlear filtering, quantified as envelope amplitude modulation, was developed to help interpret the main features of the behavioral data.

II. EXPERIMENT 1: REPETITION DETECTION WITHIN TONE CLOUDS

The aim of this experiment was to map performance across the parameter space of the tone clouds for a within-trial repetition detection task. This task has been extensively used to probe sensory memory with white noise (Guttman and Julesz, 1963; Warren *et al.*, 2001; Kaernbach, 2004; Rajendran *et al.*, 2016). Noises of different durations, or different duty cycles for the repetition, have been used to vary the informational content of stimuli in this task (Kaernbach,

2004; Warren *et al.*, 2001), with longer and hence richer noises leading to poorer performance. Here, we apply the same repetition detection task to tone clouds, keeping duration constant but varying the density of tone clouds over a broad range of densities to vary the informational content. This experiment also served as a pilot for the subsequent main experiment probing perceptual learning, part of which replicated Experiment 1 with a larger group of participants. It was also aimed at testing our initial hypothesis that denser tone clouds would induce lower performance in the within-trial repetition detection task.

A. LISTENERS

There were four listeners (1 male, 3 female) aged 23–33 years old ($M = 27$ years), each with normal pure-tone thresholds, defined as thresholds within 20 dB HL over the audiometric frequencies 500 Hz, 1 kHz, 2 kHz and 4 kHz. Three of the listeners had not previously participated in experiments involving tone clouds or noise. The fourth listener was the first author.

B. STIMULI

1. *Tone Clouds*

The tone clouds were generated as combinations of many 50-ms pure-tone pips. The onsets and frequencies of each pip were assigned according to a time/log-frequency grid (see Figure 1, panels A/C/E), such that a pip would start within each cell of the grid at a random frequency and a random point in time, without further restrictions. The grid was defined by two parameters, the number of columns of time periods per second and the number of rows of log-frequency regions per octave. These two parameters fully specified the tone clouds. In this first experiment, there were 2, 32, or 512 columns per second crossed with 1, 4, 16, 64, or 256 rows per octave to form 15 conditions. The “density” of the tone clouds can be defined numerically by multiplying these two parameters, in units of tone pips per second per octave (tones/s/oct).

----Insert Fig. 1 about here----

For the least dense tone clouds (Fig. 1A), nearly all of the tone pips would have been resolved by the auditory system. This is illustrated with a “cochleagram” of the stimulus (Fig. 1B). As the tone clouds get denser (Fig. 1C, 1E) there are many more tone pips, but many of these would still be resolved (Fig. 1D, 1F). The example of Fig. 1E is still only moderately dense (128 tones/s/oct) relative to some other tone clouds in the experiment, but some overlap between the representation of individual tones is already visible. For the densest tone clouds (131072 tones/s/oct), individual tone pips are no longer resolved by the auditory system (Fig. 1H) and the stimuli appear (and sound) similar to a spectrally matched noise (Fig. 1G).

Within each cell, the onset time, log-frequency, and phase of the pips were assigned randomly with rectangular distributions. As tone pips were 50 ms long, they may often end in a subsequent column of the grid. Since the nominal duration of the tone clouds (here 1 second) refers to the size of the grid, the actual duration of the tone cloud includes the extra time required for the last tone pips to be completed (thus up to 1.05 s in total). As such, the average number of concurrent tone pips gradually increased over the first 50 ms and decreased over the last 50 ms. The individual tone pips were themselves shaped by 50-ms long Hanning windows. The design of the tone clouds thus implicitly includes a natural fade in and fade out, stemming partly from the increasing or decreasing number of concurrent tone pips and partly from the gradual onset and offset of each tone pip.

The maximum amplitude of each tone pip was uniform over the frequency range 200–3200 Hz. Since the tone pips were distributed evenly over log frequency, there was therefore an equal amount of power in each octave. This induces a 3 dB/octave decrease when measured in terms of spectral density, making the average spectrum of the tone clouds “pink”, rather than “white”.

For two octaves at the extremes of the tone clouds (50–200 Hz and 3200–12800 Hz), the tone pips’ amplitudes were reduced by 30 dB/octave. These “fade outs” in the frequency domain were included so that denser tone clouds would not have abrupt edges in the spectral domain, which could have emphasized tone pips near the edge (Small and Daniloff, 1967; Carterette *et al.*, 1969). Within this overall spectral envelope, the peak amplitude of each tone pip was set to be inversely proportional to the square root of the number of tone pips per second. Thus, the average power of the tone clouds was independent of its parameters, but the stimuli retained small variations in total power between individual exemplars, stemming from the precise set of frequencies in a given tone cloud relative to the imposed spectral envelope and, to a lesser extent, the phase of each of the tone pips.

Repeated tone clouds (RTs) were generated by copying the pattern of onsets, frequencies, and phases from one half-second region of the time-frequency grid to the next. Each of the tone cloud schematics in Figure 1, panels A/C/E, represent such repeated stimuli. The advantage of repeating the tone clouds in the grid domain, as opposed to the acoustical domain, is that the transition from the first segment of the tone cloud to its repetition is guaranteed to be seamless, not statistically different from any other column-to-column transition. By design, this transition contains no local cues that distinguish it from an unrepeated sound: the sound at any point depends entirely on the pattern of frequencies, phases, and onset times in the preceding columns, which come from the same pseudorandom distribution whether or not the tone cloud is repeated. This cross fade between the first and second segment is, in fact, equivalent to the sum of the fade in at the start of the cloud and the fade out at the end. Thus, to be fully accurate, a repeated half-second segment of a tone cloud can be viewed as a 50-ms fade in, a 450-ms segment, a 50-ms crossfade, an identical repetition of the 450-ms segment, then the 50-ms fade out. Unrepeated tone clouds (Ts) were generated to be the same length as the RTs, but with

freshly generated onsets, frequencies, and phases throughout the grid. Example sound files as well as Matlab routines to generate tone clouds are provided as Supplemental material².

2. Spectrally Matched and White Noises

For comparison with previous studies, two conditions were included in which listeners detected repetitions of noise samples. In one condition, the spectrum of the noise was shaped to match the average spectrum of the tone clouds. In the other condition, white noise was used, for comparability with Agus *et al.* (2010).

The spectrally matched noise was generated from an inverse fast Fourier transform. First, an average spectrum for the tone cloud was calculated, by taking an RMS average of the Fourier components for 100 000 one-second tone clouds, zero padded to give a total of 2^{15} samples. This spectrum was then approximated to within 0.1 dB as a function that was piece-wise linear in terms of log frequency (in octaves) and log spectral density (in dB). As such, the spectral density at any given frequency component could be efficiently looked up or interpolated as necessary. Then, in order to generate spectrally matched noises of any duration, the required spectral density was calculated at as many frequency components as there were samples of noise. The phases were randomized and the amplitude of each component was multiplied by random Rayleigh-distributed values to reflect the usual statistics of stationary noises (Rice, 1952; Hartmann, 1998, pp 605–609). The waveform was then generated from this frequency-domain representation of the noise by an inverse FFT. Finally, the noise was ramped on and off with 25 ms raised-cosine ramps.

In order to generate a spectrally matched noise that was repeated, the odd-frequency components were set to zero. Since each of remaining even-frequency components repeated at

a multiple of 2 Hz, the noise as a whole naturally repeated twice per second without any acoustical artifact at the join.

The white noise was generated as normally distributed random values at a sample rate of 44.1 kHz. For the repeated white noise, half this number of samples were concatenated to an identical copy of themselves. Since each sample of white noise is statistically independent, the join between the first and second segment was also acoustically seamless. As for the spectrally matched noise, the white noise was ramped on and off with 25 ms raised-cosine windows. This differs from Agus *et al.* (2010), whose noises started and finished abruptly, but resembles more closely the temporal dynamics of the tone clouds used here.

C. PROCEDURE AND APPARATUS

Listeners were first trained to detect repetitions in tone clouds in a yes-no task, starting with ten repetitions of half-second snippets (and nonrepeating controls of equal duration) and gradually reducing the number of repetitions to three. Since there were many conditions, listeners were just trained on a representative subsample of the conditions: moderately dense clouds (16 columns per second $\times$ 8 rows per octave), sparse clouds (4×2), and finally very dense clouds (in fact, the spectrally matched noise was substituted here to save computational time). All listeners were trained on these three conditions in that order. This lasted approximately 45 minutes.

Each trial was triggered by a key press. After a fixed delay of 0.5 s, the stimulus was presented. Listeners responded as to whether the stimulus was repeated or not using a key press, which in turn triggered the next trial after a brief pause. Listeners received feedback on each trial throughout both the training and the main experimental blocks.

After training, listeners completed a total of 51 blocks, each including 50 RTs consisting of two half-second snippets and 50 unrepeated Ts with the same total duration, interspersed in a random order. In each block, repetition detection was measured for one of the 3×5 combinations of tone-cloud parameters or one of the 2 types of noise. Each of these 17 blocks was completed once in a random order, and this was repeated twice more with different random orderings of blocks (3×17 blocks in total). Listeners took breaks between blocks as required and were allowed to pace themselves through the 5.5 hours (on average) in 3–7 sessions per participant.

The stimuli were presented diotically at an overall level of 70 dB(A) in a double-walled sound-treated booth, through an RME Fireface UC soundcard and Sennheiser HD600 headphones.

D. ANALYSIS

Each listener's sensitivity to repetitions was measured in terms of the signal-detection theory measure, d' , calculated for each condition as the difference between z -transformed hit rates and false-alarm rates (Macmillan and Creelman, 2005, p. 8). Hit rates and false-alarm rates were averaged across blocks before z -transformations. No correction for extreme scores was made, because no participant maintained a perfect score throughout any condition.

Effects of tone-cloud parameters were tested for in a 3×5 repeated-measures ANOVA using d' as the dependent variable. A follow-up three-way, repeated-measures ANOVA tested whether there was an effect of tone cloud construction at the one density at which tone clouds were generated three different ways (512 tones/oct/s).

E. RESULTS

Figure 2 shows the average sensitivity in each condition, measured as d' , plotted against the “tone-cloud density” in tones/s/oct calculated as the multiple of the columns per second and rows per octave. Broadly, the tone-cloud density seems to account for many of the differences between conditions, with the highest repetition detection occurring for the least dense tone clouds. Sensitivity dropped with increasing density up to around 128 tones/s/oct, after which sensitivity was relatively constant around a mean $d' = 1.0$. A 3×5 repeated-measures ANOVA on all tone-cloud conditions (i.e. excluding the noise condition) showed significant effects of the number of columns per second ($F_{2,6} = 29.45$, $p = 0.001$), the number of rows per octave ($F_{4,12} = 11.33$, $p < 0.001$), and a significant interaction between the two ($F_{8,24} = 8.34$, $p < 0.001$).

Figure 2. Mean repetition-detection sensitivities observed in Experiment 1 for tone clouds with 15 different parameter sets and for the two noise stimuli. Error bars are 95% confidence intervals centered on the mean.

For the two noise stimuli, repetition detection also averaged $d' = 1.0$, the same as the denser tone clouds. This is generally higher than the repetition detection observed by Agus *et al.* (2010), perhaps reflecting the greater experience of these listeners, especially towards the end of this longer experiment. The sensitivity to repetitions in white noise ($M = 0.9$) was similar to that for spectrally matched noise ($M = 1.0$; $t_3 = 0.51$, $p = 0.64$).

For medium-dense tone clouds (64–512 tones/s/oct) there was a consistent trend of less sensitivity to repetitions in tone clouds at 32 columns per second compared to other equally dense tone clouds, but these differences were not significant. For example, at 512 tones/s/oct, where the same density of tone cloud was obtained in three different ways, a repeated-measures ANOVA showed no significant effects of tone-cloud construction ($F_{2,6} = 3.24$, $p = 0.11$).

F. DISCUSSION

Within the range of tone-cloud parameters tested, there was a wide range of mean sensitivities observed, with d 's ranging from 0.8 to 3.2. This shows that tone clouds can be parametrically controlled to vary the level of difficulty for a repetition detection task, while keeping the stimulus duration constant. The decrease in performance observed for longer white noises (Guttman and Julesz, 1963; Kaernbach, 2004) may thus reflect both the temporal decay of auditory memory and a change in complexity of the stimulus (cf. Goossens *et al.*, 2008). Note that Kaernbach (2004) also varied the time between repeated segments of constant length and observed a reduction in performance, which can only be explained in terms of memory decay. Therefore, in the repetition detection task, as in most instances of memory-constrained tasks, there is both an influence of stimulus complexity and time course on performance, beyond what has been termed echoic memory (Cowan, 1984; Nees, 2016; but see also Demany and Semal, 2008).

On a methodological level, the performance level of participants was mostly determined by the density of the tone clouds. This provides a simple handle to design relatively easy or relatively difficult repetition detection tasks by manipulating the density of the clouds, up to a point. Introspectively, manipulating the density changes the sounds from melody-like to noise-like, with a continuous transition between the two in terms of the underlying stimulus parameters. As such, the level of task difficulty could be adapted to each participant or to different experimental settings.

As an exception to this simple rule, one trend was observed suggesting a more difficult repetition detection task at medium-dense tone clouds, and not for the densest clouds or for noises. The dip in performance did not reach our statistical significance threshold, but, with only four listeners, there was little statistical power to detect a small effect, leaving some ambiguity. Therefore, two of the more distinct points that suggested this trend (at 512 tones/s/oct) were re-measured with more participants as part of the main experiment, which is presented in the next section.

III. EXPERIMENT 2: LEARNING TONE CLOUDS

This experiment aimed to apply the perceptual learning paradigm of Agus *et al.* (2010) to tone clouds of different densities. Such a paradigm has been mostly used with noise as a stimulus (Agus *et al.*, 2010; Agus and Pressnitzer, 2013; Luo *et al.*, 2013; Agus *et al.*, 2014; Andrillon *et al.*, 2015; Andrillon *et al.*, 2017). It has also been applied to stochastic click trains (Kang *et al.*, 2017) and rapid melodies (Bianco *et al.*, 2020) as well as visual (Gold *et al.*, 2014; Maharjan *et al.*, 2017; Kang *et al.*, 2018) or tactile stochastic stimuli (Kang *et al.*, 2018), both alone and in multimodal combinations (Kang *et al.*, 2018).

In all of these studies, the principle remained the same. From the perspective of the participants, the task was identical to that of our Experiment 1: a within-trial repetition detection task. However, identical copies of a specific exemplar of the stochastic stimulus re-occurred over several trials within each block. Performance improved over the course of a block for the re-occurring stimulus, which was interpreted as perceptual learning. Here we used tone clouds of varying densities as the stochastic stimuli. Their parameters were chosen, following the outcome of Experiment 1, to vary from sparse to dense clouds. The main aim was to test whether rapid learning could be observed over a broad range of stimulus complexity and baseline performance.

A. LISTENERS

There were ten listeners (1 male, 9 female) aged 21–30 years old ($M = 25$ years). All listeners had normal pure-tone thresholds, except for one listener who had slightly raised thresholds at some high frequencies in both ears (20–25 dB HL at 4 kHz). All of the listeners were naïve as to the presence of a re-occurring stimulus that could be learnt within a block. None of them had previously participated in repetition-detection experiments involving either noise or tone clouds.

B. STIMULI

The stimuli were tone clouds and a spectrally matched noise, synthesized as described in Experiment 1. As well as the repeated and unrepeated tone clouds of Experiment 1, there was a third type of stimulus, the “reference repeated tone cloud” (RefRT). RefRTs were generated in the same manner as the other repeated tone clouds (the RTs), but a given RefRT re-occurred identically on multiple trials randomly interspersed throughout an experimental block.

Building on the results of Experiment 1, there were seven different densities of tone cloud, formed as: 4 columns per second $\times$ 2 rows per octave, 8×4 , 16×8 , 32×16 , 128×64 , and 512×256 . In order to test whether the manner of construction of the tone clouds mattered, as suggested by trends in Experiment 1, half of the “standard” 32×16 blocks were replaced by “special” 2×256 blocks, which involved the same overall number of tone pips but with a different spectrotemporal distribution.

C. PROCEDURE AND APPARATUS

Listeners were first trained to detect repetitions in tone clouds and noise using the same training regimen as described for Experiment 1. Only T and RT stimuli were included in the training and trial-by-trial feedback was provided. The RefRT was only introduced during the experimental blocks, where there was no feedback. This was so that listeners were not cued towards the RefRTs. Any learning would have to be entirely unsupervised.

The experimental blocks consisted of 80 trials, with 40 T, 20 RT and 20 RefRT trials. Thus, half of the stimuli were repeated and half were unrepeated. Each block tested one tone-cloud condition at one of the six densities, or spectrally matched noise. These seven conditions formed a series of seven blocks that were repeated eight times (each time in a random order) resulting in a total of 56 experimental blocks per participant. For the tone-cloud density that could be formed in two different ways (32 columns per second $\times$ 16 rows per octave versus 2×256), each pair of series included just one block of each in a random order.

The equipment used was the same as for Experiment 1. Listeners completed the experiment in an average of 5 hours, spread over three visits.

D. ANALYSIS

As for Experiment 1, d' was calculated for each listener and condition, here calculated separately based on z -transformed hit rates for RT and RefRT. Perfect scores were mapped to the equivalent of $1/2N$ errors (Macmillan and Creelman, 2005, p. 8), where N is the 160 trials for hit rates or the 320 trials for false-alarm rates. Following the methods of Agus *et al.* (2010), these two d 's are calculated from the same false-alarm rate, as false alarms for RT and RefRT are intermingled and indistinguishable.

These d 's are reported in the text and Figures 4 and 5, for consistency with Experiment 1 and related work. However, as pointed out by a reviewer, d 's for RT and RefRT covary, to some degree, because of their shared estimates of the false-alarm rates. This violates the assumptions of the ANOVA when comparing results across conditions (but does not affect the comparison of RT and RefRT performance within conditions). As such, the statistical tests reported for this experiment are based on z -transformed hit rates alone, without the subtraction of z -transformed false-alarm rates to calculate d' . However, we also verified that using d' as the dependent variable here did not meaningfully change the outcome of the analysis.

The z -transformed hit rates were submitted to an 8×2 repeated-measures ANOVA, with factors of tone-cloud density and stimulus type (RT or RefRT). Post-hoc comparisons were based on equivalent repeated-measures ANOVAs with fewer tone cloud densities, as indicated in the description of the results below.

E. RESULTS.

1. Sensitivity

Figure 3 shows the sensitivity, measured as d' , for the repetition detection task in each of the tone-cloud conditions plus the noise stimulus (considered in the following description as a tone

cloud with infinite density). For RTs—that is, tone clouds that were presented in only one trial and did not re-occur during a block—the pattern of results replicated the findings of Experiment 1. Sensitivity decreased as density increased up to around 512 tones/s/oct. For RefRTs, sensitivity followed exactly the same pattern with respect to density. However, importantly, performance was better for RefRTs than RTs at all densities. In other words, tone clouds that re-occurred over several trials led to better within-trial repetition detection performance.

Figure 3. Mean repetition-detection sensitivities observed in Experiment 2, for tone clouds containing a single within-trial repetition but presented on only one trial (RT) and tone clouds containing a within-trial repetition that re-occurred during an experimental block (RefRT). Error bars are 95% confidence intervals centered on the mean.

The repeated-measures ANOVA showed that there was a significant difference between RT and RefRT ($F_{1,9} = 40.46, p < .001$) as well as a main effect of density ($F_{7,63} = 31.55, p < .001$), most of which presumably comes from the high performance observed for the two least dense

stimuli (average $d'_{RT} = 3.5$ and 1.9 , compared to d'_{RT} comprised between 0.6 – 1.0 for all other conditions). Interestingly, there was no interaction between density and stimulus type ($F_{7,63} = 0.18, p = .99$). This shows that the bonus in performance for RefRTs was statistically about the same across all tone-cloud densities. This was the case even though baseline difficulty varied substantially, as estimated by RT performance.

To follow up on Experiment 1, for the density of 512 tones/oct/s, two grid types were included. The standard tone cloud had 32 columns/s $\times$ 16 rows/oct, in contrast to a special tone cloud (2×256). We replicated the trend observed in Experiment 1. The standard tone cloud produced poorer performance than the special ($d'_{RT} = 0.6$ for the 32×16 grid compared to $d'_{RT} = 0.1$ for the 2×256 grid). Perhaps because of the larger number of participants, the difference was now significant as estimated by a 2×2 repeated measures ANOVA (including both RT and RefRT; $F_{1,9} = 7.66, p = .02$).

Finally, focusing on the denser standard tone clouds (16×8 , 32×16 , 128×64 , and 512×256), for which participants' performance seemed to have bottomed out, there remained an effect of stimulus density in a 4×2 repeated-measures ANOVA ($F_{3,27} = 4.42, p = 0.01$). Post hoc pairwise comparisons (without correction) suggested that this effect was driven by the 32×16 tone cloud, whose repetitions were detected significantly less than other stimuli, specifically the 16×8 stimulus ($p = 0.03$) and the more dense 512×256 stimulus ($p = 0.01$). Thus, sensitivity did not reduce monotonically with density, but rather had a small dip in performance around 512 tones/s/oct for the standard tone clouds.

2. Time course

The advantage in sensitivity for RefRTs over RTs is suggestive of perceptual learning, as sounds that re-occurred during a block induced better performance overall. Evaluating how

1 response patterns changed over the time course of a block provides an opportunity to see how
2 performance changes. Indeed, initially RefRTs and RTs are drawn from the same random
3 process, so they should produce identical response patterns for the first trial of a block on
4 average. In particular, an analysis of the dynamics also indicates the rapidity of any learning,
5 which has been shown to vary across stimuli and modalities (Kang *et al.*, 2017; Kang *et al.*,
6 2018).

7 The time-courses were analyzed by the method described in Agus *et al.* (2010) and Kang *et al.*
8 (2017). Hit rates for RefRTs and RTs were averaged across participants for each presentation
9 of the RefRT (identical throughout a block) and RTs (different on each trial throughout a
10 block). Figure 4 shows these hit rates as a function of position in a block: for the first
11 presentation, second, and so on. Generally, performance was initially similar for RefRTs and
12 RTs, but after a few trials, the responses to the two types of stimuli diverge, with better
13 performance for RefRT than RT. One of two types of dynamic response was observed for 7
14 out of the 8 conditions: either there were more hits for RefRTs toward the ends of the blocks
15 or there were fewer hits for RTs toward the end of the blocks. As shown by Agus *et al.* (2010),
16 this drop in RT hit rate can reflect a change in criterion: as RefRTs become easier to detect, the
17 criterion for repetition detection becomes more conservative, thus leading to fewer hits for
18 RTs. Under this interpretation, both patterns of performance signal that there was learning. No
19 learning, by contrast, would be evidenced by no change in performance in either condition over
20 the course of a block.

21 We fitted flat line and exponential models to the average data, plotting that which best
22 explained the data, taking into account the added complexity of the exponential model
23 compared to the flat line model for the best model selection (Agus *et al.*, 2010). Results of the
24 analysis are summarized in Table 1. In general, we observed either a significant increase in hit

rate for RefRTs for the densest clouds or a significant decrease in hit rate for RTs in the sparsest clouds. As is clear from Figure 3, hit rates were close to ceiling for the sparsest clouds even for RT stimuli, which could explain why RefRT performance did not increase over the course of a block. Rather, as RefRT trials became even more obvious when the block progressed, the criterion became more conservative. There is one exception to this description for one middle density (32×16 tones/oct/s) where there were no significant changes for either RTs or RefRTs.

<i>Tone-cloud density</i> (tones/s/oct.)	<i>RT</i>			<i>RefRT</i>		
	<i>F</i>	<i>p</i>	half-life	<i>F</i>	<i>p</i>	half-life
4×2	9.21	.002	2.4	0.99	ns	—
8×4	6.26	.009	5.4	0.75	ns	—
16×8	11.46	.001	1.8	2.50	ns	—
32×16	0.39	ns	—	1.09	ns	—
2×256	1.56	ns	—	7.18	.005	18.0
128×64	1.95	ns	—	8.84	.002	1.8
512×256	0.98	ns	—	9.76	.002	0.79
matched noise	10.22	.001	7.3	1.09	ns	—

Table 1. Statistics for the time course of hit rates for RT and RefRT stimuli over trials (see also Figure 3). Where there were significant changes in hit rate over the series of trials, estimates of the half-life of the exponential fit are also listed.

1

2

3

4

5

6

7

8

9

10

Figure 4. Time courses of the hit-rate across blocks for tone clouds of different densities and constructions (A–G) and spectrally matched noise (H) separated into RTs (black circles) and RefRTs (red circles). The x -axis represents the presentation order within each type of stimulus across a block, not counting the intervening trials of other stimulus types. The fitted lines are exponential where the hit-rate was observed to change significantly, or straight averages otherwise (see text for details), and the shaded areas around these indicate asymptotic 95% confidence intervals centered on the mean. The hit rates on the first presentation of each stimulus type were expected to be the same, whereas any subsequent difference between RT and RefRT is indicative that re-occurrence induced learning of the RefRTs.

Nonetheless, the pattern of results in this case displays identical performance initially and an advantage for RefRT later on in the block. The failure to reach significance could be due to statistical noise in the data, in particular with inexplicable dips in performance in the second half of the block. Note that for the equally dense 32×16 and 2×256 conditions, there were half the number of blocks relative to each of the other conditions. In summary, all conditions but one showed evidence of temporal changes in performance consistent with rapid learning, with the exception likely related to statistical noise.

The rate of learning was estimated by a fitted “half-life” parameter representing the number of presentations of the stimulus over which the hit rate gets closer to its asymptote by half. Where there were significant changes, the fitted half-lives of learning were generally of the order of a small number of trials, in line with the rapid learning observed previously (Agus *et al.*, 2010; Kang *et al.*, 2017). The longer half-lives were not particularly reliably fitted: their estimated 95% confidence intervals included the possibility that the half-life was a single trial.

F. DISCUSSION

As had already been shown with noise, click trains and tones sequences, we observed here an advantage in within-trial repetition detection performance for stochastic tone clouds that re-occurred throughout an experimental block, compared to tone clouds that were only ever presented in a single trial. This advantage was observed for all densities tested, which included vastly different perceptual regimes: essentially, from melody-like stimuli to noise. Moreover, an accompanying rapid change in response patterns over time was observed for most densities. Together, these findings suggest that rapid auditory learning applies to stochastic spectrotemporal patterns, from simple melody-like patterns to more noise-like tone clouds.

Learning was observed for all types of stimuli despite large effects in the baseline repetition detection rates, with d' 's for RT ranging 0.6–3.5. Perhaps surprisingly, the amount of learning measured as the difference in d' sensitivity to RT and RefRT was approximately constant, with a benefit of 0.4–0.6 d' units for RefRT. Interpreted literally in the signal detection theory framework, this would suggest that learning shifted the perception of the RefRT stimuli by a constant amount in terms of an internal perceptual dimension normalized with respect to the variability in the stimuli (cf. Macmillan and Creelman, 2005, pp. 16-18). However, the size of the learning effect could have been impeded by ceiling effects for the easiest conditions. Nevertheless, such a robustness of the effect over a range of conditions suggests that it is not necessary for future studies to aim precisely for a particular baseline task difficulty in order to probe rapid auditory learning.

A secondary aspect of this experiment was to revisit a comparison from Experiment 1, testing whether performance was explained by tone-cloud density alone, or whether the manner of construction also affected performance (specifically, whether the exact split into columns per second and rows per octave matters). This appears to be the case, to a small extent: a difference in d' of 0.35. Thus, although tone-cloud density largely predicts repetition-detection abilities, there are also smaller effects related to the construction parameters of the tone clouds. When the temporal position of the tone pips was not as carefully controlled as in the standard generation procedure, repetitions were detected more easily (based on a comparison of 2×256 and 32×16 conditions). Qualitatively, less control over the timing of the tone pips could lead to chance clustering of tone pips and, conversely, regions of sparsity. Then, these denser and sparser regions could provide salient cues for repetition detection. This qualitative interpretation is supported by a functional model, described in the next section.

IV. A MODEL OF TONE-CLOUD AMPLITUDE MODULATION CUES

Some of the differences in baseline performance for different constructions of tone clouds qualitatively suggested that repetition detection was affected by the evenness of the spread of tone pips across the tone clouds. The model presented here aims to explore potential cues further by measuring how evenly power is spread spectrotemporally, as quantified by measuring the depth of amplitude modulation observed after simulated cochlear filtering.

Behaviorally, in both experiments, the ability to detect repetition in RT stimuli varied widely with respect to the tone density, with sparser tone clouds providing generally better performance than denser clouds (Figs. 2, 3). If the repetition detection task is taken as a measure of sensory memory capacity (Kaernbach, 2004), an intuitive and appealing explanation for this finding could be that performance depended on the number of tones to be stored in memory, with more tones lowering performance. However, such an interpretation is not consistent with finer aspects of the results. Performance generally decreased from sparse to denser clouds, albeit slightly worse for medium tone-cloud densities than for the densest clouds. In particular, different performance levels were observed for tone clouds with the same density: a 32×16 grid produced worse performance than a 2×256 grid, even though the average number of tones per stimulus was the same in the two cases (Fig. 3). We thus hypothesized that this latter observation could be related to the spread of tone pips across the tone clouds.

To quantify this qualitative hypothesis, we probed a functional model of the amplitude modulations available to participants for the different tone-cloud stimuli of Experiment 2. The depths of amplitude modulations indicate the extent to which the stimuli are evenly distributed in terms of spectrotemporal density over time. Briefly, the model estimated amplitude modulations (AM) within simulated cochlear channels as they respond to RTs. A gammatone filterbank was used to model peripheral frequency selectivity (Patterson *et al.*, 1995).

Peripheral filter widths impact the AM cues that are transmitted to subsequent processing stages, as AM frequencies higher than half a peripheral filter width would be highly attenuated. We then computed a modulation spectrum within cochlear channels, to estimate the power at various AM frequencies. Similar approaches have been used before (e.g., Dau *et al.*, 1997), in particular to estimate speech intelligibility (Steeneken and Houtgast, 1980; Elhilali *et al.*, 2003; van Wijngaarden and Houtgast, 2004). We sampled different peripheral filters and AM frequencies, to test whether some combinations of the parameters would predict the pattern of performance observed in Experiment 2. Specifically, we hypothesized that higher modulation power should be associated to more salient acoustic events and better performance at detecting repetitions.

A. METHODS

To calculate modulation spectra, a twice-repeated stimulus was first filtered to represent the transfer properties of the middle ear, using a 1st-order Butterworth bandpass filter between 400 and 8500 Hz. The output of the transfer filter was then fed to gammatone filters that were 1-ERB wide (Patterson *et al.*, 1995; Slaney, 1998) at 49 center frequencies, six per octave distributed logarithmically from 50 Hz to 12 kHz. The envelope of each filtered signal was then extracted by half-wave rectification and lowpass filtering at half the center frequency of the gammatone filter. The modulation spectrum was finally calculated as the fast Fourier transform of the second half-second of this stimulus, with modulation depth converted to decibels, relative to full-wave modulation. On this scale, 0 dB is equivalent to 100% modulation between silence and maximum amplitude and a constant pure tone would have a modulation depth of $-\infty$ dB. An implementation of the full model is provided as Supplemental material³.

Note that all of the filters used were feed-forward filters, with each sample of the envelope influenced by relatively few preceding samples, so the second segment of a twice-repeated noise is the same as a half-second portion taken from an indefinitely repeated version of the same segment. As such, this half second envelope can be treated as a periodic signal, and no further windowing was necessary or desirable before calculating its modulation spectrum.

The types of signals generated were the same as the stimuli generated for Experiment 2, that is, tone clouds of different densities and spectrally matched noises. For each analysis, 1000 exemplars of each stimulus type were generated. These were generated afresh, as the exact exemplars used in Experiments 1 and 2 were not stored. The average modulation spectra were calculated as root-mean-squares of the amplitudes at each modulation frequency component. As modulation depth is here used as a proxy for a wide range of features that may be exposed in local peaks and dips, modulation phase was not considered, even though phases may also be important, particularly at lower modulation rates (e.g., Dau *et al.*, 1996).

To indicate how well listeners' variations in repetition detection performance could be predicted by modulations at any given center frequency and modulation rate, Pearson's r was calculated from the mean modulation depths for different types of stimuli and the mean performance for RT stimuli in Experiment 2.

B. RESULTS

The top panel of Figure 5 shows one example of the results obtained by the model, specifically for a modulation frequency of 4 Hz in an auditory channel centered at 400 Hz. For this example, the amplitude modulation power was closely correlated with mean repetition detection in the RT stimuli in the eight conditions (shown in light gray; $r(6) = 0.97$, $p < 0.001$) and reflects the key features of the behavioral data of Figure 3:

Figure 5. A: Modulation depths at 4 Hz at a center frequency of 400 Hz for the tone clouds and noise presented in Experiment 2, including “standard” tone clouds, whose rows and columns double proportionately, and “special” tone clouds with just two columns and many rows (see text for details). The gray data and corresponding gray axis reproduce the mean repetition-detection sensitivities of RT from Experiment 2, originally shown in Figure 3. **B:** Heat-map showing the strength of correlation between the mean RT repetition detection measured behaviorally (see panel A) and average modulation depths obtained from the model over a range of center frequencies and modulation frequencies. All positive correlations are shown, from dark to bright. **C:** As for panel B, but with correlations based on only six points, omitting the two least-dense tone-cloud conditions.

- 1 1. For sparser tone clouds (< 512 tones/oct/s), modulation power decreases rapidly with
2 increasing density.
- 3 2. There was more modulation power for the 2×256 tone-cloud relative to the 32×16
4 tone cloud.
- 5 3. There was *less* modulation power for the standard tone cloud at 512 tones/oct/s relative
6 to denser tone clouds and noise.

7 Thus, even the finer and unexpected aspects of the behavioral results, such as the dip in
8 performance around 512 tones/s/oct, are mirrored in the model simulations of amplitude
9 modulation power.

10 The model strongly correlated with the behavioral data over a wide range of center frequencies
11 (see Figure 5B) up to modulation rates around 35 Hz, depending on the center frequency. These
12 correlations were likely strongly driven by the two least dense tone clouds, which had the
13 deepest modulations and the most reliably detected repetitions (see point 1 above).

14 To focus on points 2 and 3 above, we repeated the analysis with these two tone least dense tone
15 clouds excluded from the analysis (see Figure 5C). There remained strong correlations between
16 the model and behavioral data, at slower modulation rates (2–6 Hz) across a wide range of
17 center frequencies. Correlations between behavioral data and the model for slow amplitude
18 modulations are consistent with the idea that some tone-cloud designs spread energy more
19 evenly over time and frequency than others. In this focused version of the model analysis,
20 correlations were also observed in other regions of the parameter space, for high modulation
21 rates in particular (40-50 Hz).

C. DISCUSSION

Although the sensitivity to repetitions in tone clouds could broadly be predicted by modulation depth for a wide range of modulation rates at a wide range of center frequencies, the more surprising aspects of the data are only also predicted by the slowest modulation rates (2–6 Hz). To put these parameters in context, acoustic analyses of speech and music databases (Ding *et al.*, 2017; Poeppel and Assaneo, 2020) have observed broad spectral peaks in the same range (5 Hz for speech, 2 Hz for music). Perceptually, we are also most sensitive to these relatively low modulation frequencies (Viemeister, 1979). Finally, for speech intelligibility, envelope modulations around 4 Hz seem critical (Drullman *et al.*, 1994). Fast modulation cues were also correlated to behavior when the sparsest tone clouds were excluded from the analysis (Fig. 5C). The range for both fast and slow modulations is broadly consistent with the asymmetric sampling model developed for speech (Giroud *et al.*, 2020), which has previously been implicated in the neural bases of the memory for noise (Luo *et al.*, 2013). However, there is a large variability in Figure 5C because each square is based on the correlation of only six points with little variation within these points (cf. the six right-most points of the behavioral data of Figure 5A), so the interpretation of the high modulation frequency cues must remain tentative.

In any case, the strong correlation between the behavioral results and the modulation model over a wide range of parameters suggests that whatever features the listeners are using to detect the repetitions, they co-vary with the modulation spectrum and could even include features based in the modulation domain. For the sparsest tone clouds, it is easy to imagine that bunching of tone pips would lead to clearer spectrotemporal peaks or dips that would aid repetition detection, particularly if different modulations are correlated (e.g., McDermott and Simoncelli, 2011). This could also account for the improved repetition detection for the $2 \times$

256 tone cloud, whose tone pips were less constrained in time and freer to clump than the equally dense 32×16 condition.

From this perspective, the correlation of modulations and behavior would not necessarily show that listeners are learning the temporal envelope of sounds; rather, the correlation could merely reflect that the depth of modulations covaries with the presence of other perceptually important features, whatever they are. For the denser tone clouds and noise, a model based on slow modulations also does not seem immediately consistent with previous experimental data that suggests that noise repetition detection or recognition is based on narrower temporal windows, at most 100 ms (Kaernbach, 1993; Agus *et al.*, 2010) which can be more or less localized in frequency (Kaernbach, 1993). Perhaps deeper slow modulations allow the listeners to focus their attention on shorter windows of the stimulus, which has been shown to aid listeners in noise repetition detection or recognition (Warren *et al.*, 2001; Goossens *et al.*, 2008). Alternatively, dips in power could expose sparser regions of moderately dense tone clouds, in which smaller numbers of tone pips could be audible, with their more distinctive, less noise-like features.

A finer point relates to the 32×16 tone cloud, which produced the lowest modulation depth overall at low modulation rates. It is perhaps surprising to stumble across a non-stationary sound that is complex and constantly changing, yet in some senses smoother than noise in terms of its modulations. This is reminiscent of low-noise noise, a broadband signal whose power is unusually stable over time, with far less modulation than would be found in a naturally occurring noise (Pumplin, 1985). However, after bandpass filtering, such as that which occurs in cochlear filters, natural fluctuations re-emerge (Hartmann and Pumplin, 1988). Here, we have demonstrated that dense tone clouds can be constructed in such a way that they fluctuate less than noise, indicating that broadband cochlear lower-noise noise is possible to an extent,

at least for slower modulation rates (here up to around 10 Hz for mid-range frequencies). The link between modulation depth and repetition detection could be tested by exploring the parameter space for even smoother tone clouds. The existence of broadband stimuli with less modulation depth than noise after cochlear filtering could be of relevance to psychoacoustical experiments that use low-noise noise (e.g., Hopkins *et al.*, 2010).

V. GENERAL DISCUSSION

The present results show that adult normal-hearing listeners can rapidly learn tone clouds over a broad range of spectrotemporal complexity, from melody-like, to texture-like, to noise-like stimuli. This class of sounds can thus be added to noise (Agus *et al.*, 2010), clicks (Kang *et al.*, 2017), and sequences of tones (Bianco *et al.*, 2020), as sounds that are amenable to rapid perceptual learning. This is an important generalization of the operating range of the type of learning mechanisms originally demonstrated with noise (Agus *et al.*, 2010) to a stimulus class arguably closer in overall statistics to natural sounds (Młynarski and McDermott, 2018). Moreover, tone clouds are a particularly useful methodological addition because their level of complexity can be tailored parametrically by varying their density. This makes them suitable as stimuli in situations where the baseline performance might need to be adjusted, for instance if there are background sounds such as MRI scanner noise (as already demonstrated in an MRI scanner; Kumar *et al.*, 2014), if training time is limited, or if working with populations that might be put off by an overly challenging task.

A. TONE CLOUDS AND SIMILAR STIMULI

The tone clouds we designed share similarities with stimuli used in several previous studies for which learning was not the main concern. Series of multi-tone complexes have often been used to investigate the effects of uncertainty in masking (Kidd *et al.*, 1994; Kidd *et al.*, 2003; Neff and Green, 1987). For these, the pure tone components occurred rhythmically, at a fixed rate,

with simultaneous onsets across all frequency bands. Pure tones with asynchronous onsets have also been used as maskers (Gutschalk *et al.*, 2008) at a cost of controlling the number of components in any given frequency band. The grid system of our tone clouds combines features of both these stimuli, controlling the frequency statistics of the components while incorporating asynchronous onsets. Such a grid system was found to have perceptual consequences, evidenced by the small but significant differences in performance for the 32×16 and 2×256 tone clouds in Experiment 2 and paralleled by the acoustical differences in their modulation spectra. Such a grid system may thus provide an additional and meaningful control of the finer details of tone clouds.

The method of generation we adopted closely resembles synthesis techniques such as the inverse short-term Fourier transform or wavelet synthesis, particularly for the densest tone clouds, in which the stimuli sound noise-like. The key differences are that the short-term Fourier transforms are based on a linear frequency scale, with onsets constrained by the window size and overlap parameters. Log frequency, as used here, is closer to more perceptually relevant frequency spacings such as musical pitch or the equal-ERB scale (Moore and Glasberg, 1983). Wavelet synthesis methods (Kronland-Martinet, 1988; Mallat, 1989) often use grids with asynchronous onsets, but the grid cells' durations usually decrease with frequency. Essentially, there may be no perfect solution in terms of varying duration, amplitude, and quantity of components with frequency when generating tone clouds. The current strategy uses a conceptually simple grid system to control the statistics of the tone cloud.

Intuitively, the tone clouds learnt here span the broadest possible range of complexity, from sparse melody-like stimuli to dense noise-like stimuli. This range straddles most everyday sounds that are typically more complex than sine waves, but with more temporal variation than

noise (cf. Samson *et al.*, 2011). In another sense of the word “complexity”, tone clouds may be used to vary the amount of temporal modulation (Ding *et al.*, 2017), but would need to be adapted in order to control for modulation rates and depths more finely. Another characteristic of the tone clouds is that, for unrepeatable tone clouds, modulations were uncorrelated over frequency and time. This is not what is observed with many natural sounds such as textures (McDermott and Simoncelli, 2011). It is not clear how the perceptual learning mechanism studied here would interact with more structured sounds, but the simplest assumption would be that more structure would favor learning. This remains to be tested, however.

B. A GENERIC MECHANISM FOR RAPID AUDITORY LEARNING?

Despite the large variation in spectrotemporal complexity of the stimuli, and the corresponding variation in baseline difficulty of the within-trial repetition detection task, the additional benefit of learning was surprisingly consistent: 0.4–0.6 when measured as d' differences, for all tested tone-cloud densities. This was observed even though the stimulus range covered melody-like, texture-like, and noise-like sounds. Moreover, such a d' difference is remarkably similar to what was observed, using the same experimental paradigm, with very different kinds of stimuli: white noise (0.5; Agus *et al.*, 2010) or click trains (0.2-1.0; Kang *et al.*, 2017). When the paradigm was extended to other perceptual modalities, vibrotactile and visual perception, the d' difference appeared to be a little less (0.3 for each; Kang *et al.*, 2018), but this remains similar and could partially reflect that learning was slower for these stimuli.

What could explain such striking similarities in the rapid perceptual learning of so many different types of sounds? The simplest hypothesis is that the auditory learning was subserved by the same underlying plasticity mechanism, in the same brain regions for all classes of stimuli tested so far. Tone clouds encompass a broad range of spectrotemporal features at different scales, so selectivity to such features could overlap with the encoding of other types of sounds.

1 Interestingly, in a high-resolution fMRI study of natural-sound encoding in humans, it has been
2 suggested that the auditory cortex expresses selectivity for a broad range of spectrotemporal
3 features (Santoro *et al.*, 2014; see also Griffiths and Warren, 2002).

4 Directly relevant for the present discussion, there is also evidence for the implication of
5 secondary auditory cortical areas in the rapid perceptual learning of some of the sounds probed
6 behaviorally. Using similar paradigms but with noise as stimuli, magnetoencephalography
7 (MEG; Luo *et al.*, 2013) and electroencephalography (EEG; Andrillon *et al.*, 2015) studies
8 concluded that the perceptual learning of noise was correlated with early neural responses
9 likely originating from secondary auditory regions. Using the more spatially precise fMRI
10 imaging technique, and with tone clouds fixed at a single level of complexity as stimuli,
11 Kumar/Bonnici *et al.* (2014) found that tone-cloud learning recruited a secondary auditory area,
12 the planum temporale, together with non-auditory regions such as the hippocampus.

13 All of this would be broadly consistent with the involvement of secondary auditory cortical
14 regions in the rapid perceptual learning of all of these classes of acoustic stimuli studied so far.
15 Note that this does not rule out the additional contribution of non-auditory areas, such as the
16 hippocampus or even frontal areas, but for those, there is little reason to suggest differential
17 treatment for the different classes of stimuli, which is then consistent with the common locus
18 hypothesis.

19 However, an additional set of findings to be accommodated is that a similar type of learning
20 was observed for vision (Gold *et al.*, 2014; Kang *et al.*, 2018) and touch (Bale *et al.*, 2017;
21 Kang *et al.*, 2018). It is unlikely that all of these stimuli are learnt in secondary auditory cortex.
22 Rather, this may require an extension to the common locus hypothesis: the plasticity

mechanisms recruited by rapid perceptual learning may be generic and implemented with similar characteristics in different brain regions.

VI. CONCLUSIONS

Learning was observed for tone clouds over a broad range of spectrotemporal complexities, ranging from sparse tone clouds, in which each individual tone pip would be audible, to extremely dense tone clouds, which resemble stationary noise. Remarkably, learning as indexed by the performance measure d' , was similar for all kinds of stimuli tested. This finding generalizes a result, originally observed as learning of noise, to a wider range of complex, broadband stimuli, from sparse melody-like sounds to dense noise-like sounds. Therefore, we suggest that rapid perceptual learning could be a meaningful process in realistic situations, to accumulate knowledge about sounds in an unsupervised and even implicit manner. Furthermore, the present study provides a spectrotemporally complex stimulus that can be parametrically manipulated to probe such a process in a flexible way, expanding the range of experimental situations amenable to rapid perceptual learning paradigms.

ACKNOWLEDGMENTS

This work was funded by the grant ANR-17-EURE-0017 (DP).

TEXTUAL FOOTNOTES

¹ Complexity as used intuitively here relates to the information content of sounds, which could be defined as the number of acoustic events per unit time (Pollack, 1968), but may also be related to temporal (Ding *et al.*, 2017) or spectrotemporal (McDermott and Simoncelli, 2011) structure.

² See supplementary material at [URL to be inserted by AIP; *SuppPub1.zip*] to listen to examples of tone clouds at various densities, both repeated (RT) and unrepeated (T). The MATLAB functions to generate more tone clouds are available at [URL to be inserted by AIP; *SuppPub2.zip*].

³ See supplementary material at [URL to be inserted by AIP; *SuppPub2.zip*] for MATLAB functions that run the model from tone-cloud generation to the analyses shown in Figure 5.

1 REFERENCES

- 2 Agus, T. R., Carrión-Castillo, A., Pressnitzer, D., and Ramus, F. (2014). "Perceptual learning
3 of acoustic noise by individuals with dyslexia," *J Speech Lang and Hearing Res* **57**, 1069-
4 1077.
- 5 Agus, T. R., and Pressnitzer, D. (2013). "The detection of repetitions in noise before and after
6 perceptual learning," *J Acoust Soc Am* **134**, 464-473.
- 7 Agus, T. R., Thorpe, S. J., and Pressnitzer, D. (2010). "Rapid formation of robust auditory
8 memories: Insights from noise," *Neuron* **66**, 610-618.
- 9 Andrillon, T., Kouider, S., Agus, T., and Pressnitzer, D. (2015). "Perceptual learning of
10 acoustic noise generates memory-evoked potentials," *Curr Biol* **25**, 2823-2829.
- 11 Andrillon, T., Pressnitzer, D., Leger, D., and Kouider, S. (2017). "Formation and suppression
12 of acoustic memories during human sleep," *Nat Commun* **8**, 179.
- 13 Bale, M. R., Bitzidou, M., Pitas, A., Brebner, L. S., Khazim, L., Anagnou, S. T., Stevenson, C.
14 D., and Maravall, M. (2017). "Learning and recognition of tactile temporal sequences by
15 mice and humans," *Elife* **6**.
- 16 Bianco, R., Harrison, P. M., Hu, M., Bolger, C., Picken, S., Pearce, M. T., and Chait, M. (2020).
17 "Long-term implicit memory for sequential auditory patterns in humans," *Elife* **9**.
- 18 Carterette, E. C., Friedman, M. P., and Lovell, J. D. (1969). "Mach bands in hearing," *J Acoust*
19 *Soc Am* **45**, 986-998.
- 20 Cowan, N. (1984). "On short and long auditory stores," *Psychol Bull* **96**, 341-370.
- 21 Dau, T., Kollmeier, B., and Kohlrausch, A. (1997). "Modeling auditory processing of
22 amplitude modulation. I. Detection and masking with narrow-band carriers," *J Acoust*
23 *Soc Am* **102**, 2892-2905.

- 1 Dau, T., Puschel, D., and Kohlrausch, A. **(1996)**. "A quantitative model of the "effective" signal
- 2 processing in the auditory system. II. Simulations and measurements," *J Acoust Soc Am*
- 3 **99**, 3623-3631.
- 4 Demany, L., and Semal, C. **(2008)**. "The role of memory in auditory perception," in *Auditory*
- 5 *Perception of Sound Sources*, edited by W. A. Yost, A. N. Popper, and R. R. Fay
- 6 (Springer Verlag, New York), pp. 77-113.
- 7 Ding, N., Patel, A. D., Chen, L., Butler, H., Luo, C., and Poeppel, D. **(2017)**. "Temporal
- 8 modulations in speech and music," *Neurosci Biobehav Rev* **81**, 181-187.
- 9 Drullman, R., Festen, J. M., and Plomp, R. **(1994)**. "Effect of temporal envelope smearing on
- 10 speech reception," *J Acoust Soc Am* **95**, 1053-1064.
- 11 Elhilali, M., Chi, T., and Shamma, S. A. **(2003)**. "A spectro-temporal modulation index (STMI)
- 12 for assessment of speech intelligibility," *Speech Communication*, 331-348.
- 13 Giroud, J., Trebuchon, A., Schon, D., Marquis, P., Liegeois-Chauvel, C., Poeppel, D., and
- 14 Morillon, B. **(2020)**. "Asymmetric sampling in human auditory cortex reveals spectral
- 15 processing hierarchy," *PLoS Biol* **18**, e3000207.
- 16 Gold, J. M., Aizenman, A., Bond, S. M., and Sekuler, R. **(2014)**. "Memory and incidental
- 17 learning for visual frozen noise sequences," *Vision Res* **99**, 19-36.
- 18 Goossens, T., van de Par, S., and Kohlrausch, A. **(2008)**. "On the ability to discriminate
- 19 Gaussian-noise tokens or random tone-burst complexes," *J Acoust Soc Am* **124**, 2251-
- 20 2262.
- 21 Griffiths, T. D., and Warren, J. D. **(2002)**. "The planum temporale as a computational hub,"
- 22 *Trends Neurosci* **25**, 348-353.
- 23 Gutschalk, A., Micheyl, C., and Oxenham, A. J. **(2008)**. "Neural correlates of auditory
- 24 perceptual awareness under informational masking," *PLoS Biol* **6**, e138.

- 1 Guttman, N., and Julesz, B. (1963). "Lower limits of auditory periodicity analysis," J Acoust
2 Soc Am **35**, 610.
- 3 Hartmann, W. M. (1998). *Signals, Sound, and Sensation* (Springer-Verlag, New York).
- 4 Hartmann, W. M., and Pumplin, J. (1988). "Noise power fluctuations and the masking of sine
5 signals," J Acoust Soc Am **83**, 2277-2289.
- 6 Hopkins, K., Moore, B. C. J., and Stone, M. A. (2010). "The effects of the addition of low-
7 level, low-noise noise on the intelligibility of sentences processed to remove temporal
8 envelope information," J Acoust Soc Am **128**, 2150-2161.
- 9 Kaernbach, C. (1993). "Temporal and spectral basis of the features perceived in repeated
10 noise," J Acoust Soc Am **94**, 91-97.
- 11 Kaernbach, C. (2004). "The memory of noise," Exp Psychol **51**, 240-248.
- 12 Kang, H., Agus, T. R., and Pressnitzer, D. (2017). "Auditory memory for random time
13 patterns," J Acoust Soc Am **142**, 2219.
- 14 Kang, H., Lancelin, D., and Pressnitzer, D. (2018). "Memory for Random Time Patterns in
15 Audition, Touch, and Vision," Neuroscience **389**, 118-132.
- 16 Kidd, G., Jr., Mason, C. R., Deliwala, P. S., Woods, W. S., and Colburn, H. S. (1994).
17 "Reducing informational masking by sound segregation," J Acoust Soc Am **95**, 3475-
18 3480.
- 19 Kidd, G., Jr., Mason, C. R., and Richards, V. M. (2003). "Multiple bursts, multiple looks, and
20 stream coherence in the release from informational masking," J Acoust Soc Am **114**,
21 2835-2845.
- 22 Kronland-Martinet, R. (1988). "The wavelet transform for analysis, synthesis and processing
23 of speech and music sounds," Computer Music Journal **12**, 11-20.

- 1 Kumar, S., Bonnici, H. M., Teki, S., Agus, T. R., Pressnitzer, D., Maguire, E. A., and Griffiths,
2 T. D. (2014). "Representations of specific acoustic patterns in the auditory cortex and
3 hippocampus," *Proc Biol Sci* **281**, 20141000.
- 4 Lim, S. J., and Holt, L. L. (2011). "Learning foreign sounds in an alien world: videogame
5 training improves non-native speech categorization," *Cogn Sci* **35**, 1390-1405.
- 6 Luo, H., Tian, X., Song, K., Zhou, K., and Poeppel, D. (2013). "Neural response phase tracks
7 how listeners learn new acoustic representations," *Curr Biol* **23**, 968-974.
- 8 Macmillan, N. A., and Creelman, C. D. (2005). *Detection theory: A user's guide (2nd edition)*
9 (Lawrence Erlbaum Associates, Mahwah, NJ).
- 10 Maharjan, S., Gold, J. M., and Sekuler, R. (2017). "Memory and learning for visual signals in
11 time and space," *Atten Percept Psychophys* **79**, 1107-1122.
- 12 Mallat, S. G. (1989). "A theory for multiresolution signal decomposition: The wavelet
13 representation," *IEEE Trans. Pattern Anal. Mach. Intell.* **11**, 674-693.
- 14 McDermott, J. H., and Simoncelli, E. P. (2011). "Sound texture perception via statistics of the
15 auditory periphery: evidence from sound synthesis," *Neuron* **71**, 926-940.
- 16 Mishra, S. K., and Panda, M. R. (2016). "Rapid auditory learning of temporal gap detection,"
17 *J Acoust Soc Am* **140**, EL50.
- 18 Młynarski, W., and McDermott, J. H. (2018). "Learning Midlevel Auditory Codes from
19 Natural Sound Statistics," *Neural Comput* **30**, 631-669.
- 20 Moore, B. C., and Glasberg, B. R. (1983). "Suggested formulae for calculating auditory-filter
21 bandwidths and excitation patterns," *J Acoust Soc Am* **74**, 750-753.
- 22 Nees, M. A. (2016). "Have We Forgotten Auditory Sensory Memory? Retention Intervals in
23 Studies of Nonverbal Auditory Working Memory," *Front Psychol* **7**, 1892.
- 24 Neff, D. L., and Green, D. M. (1987). "Masking produced by spectral uncertainty with
25 multicomponent maskers," *Percept Psychophys* **41**, 409-415.

- 1 Patterson, R. D., Allerhand, M. H., and Giguere, C. (1995). "Time-domain modeling of
2 peripheral auditory processing: a modular architecture and a software platform," J Acoust
3 Soc Am **98**, 1890-1894.
- 4 Poeppel, D., and Assaneo, M. F. (2020). "Speech rhythms and their neural foundations," Nat
5 Rev Neurosci **21**, 322-334.
- 6 Pollack, I. (1968). "Periodicity discrimination for auditory pulse trains," J Acoust Soc Am **43**,
7 1113-1119.
- 8 Pumplin, J. (1985). "Low-noise noise," J Acous Soc Am **78**, 100-104.
- 9 Rajendran, V. G., Harper, N. S., Abdel-Latif, K. H., and Schnupp, J. W. (2016). "Rhythm
10 Facilitates the Detection of Repeating Sound Patterns," Front Neurosci **10**, 9.
- 11 Rice, S. O. (1952). *Mathematical analysis of random noise* (Bell Telephone Labs, New York).
- 12 Samson, F., Zeffiro, T. A., Toussaint, A., and Belin, P. (2011). "Stimulus complexity and
13 categorical effects in human auditory cortex: an activation likelihood estimation meta-
14 analysis," Front Psychol **1**, 1-23.
- 15 Santoro, R., Moerel, M., De Martino, F., Goebel, R., Ugurbil, K., Yacoub, E., and Formisano,
16 E. (2014). "Encoding of natural sounds at multiple spectral and temporal resolutions in
17 the human auditory cortex," PLoS Comput Biol **10**, e1003412.
- 18 Slaney, M. (1998). "Auditory Toolbox, Version 2, Technical Report No. 1998-010," (Interval
19 Research Corporation).
- 20 Small, A. M., Jr., and Daniloﬀ, R. G. (1967). "Pitch of noise bands," J Acoust Soc Am **41**, 506-
21 512.
- 22 Song, K., and Luo, H. (2017). "Temporal Organization of Sound Information in Auditory
23 Memory," Front Psychol **8**, 999.
- 24 Steeneken, H. J., and Houtgast, T. (1980). "A physical method for measuring speech-
25 transmission quality," J Acoust Soc Am **67**, 318-326.

- 1 van Wijngaarden, S. J., and Houtgast, T. (2004). "Effect of talker and speaking style on the
2 speech transmission index," J Acoust Soc Am **115**, 38-41.
- 3 Viemeister, N. F. (1979). "Temporal modulation transfer functions based upon modulation
4 thresholds," J Acoust Soc Am **66**, 1364-1380.
- 5 Viswanathan, J., Remy, F., Bacon-Mace, N., and Thorpe, S. J. (2016). "Long Term Memory
6 for Noise: Evidence of Robust Encoding of Very Short Temporal Acoustic Patterns,"
7 Front Neurosci **10**, 490.
- 8 Warren, R. M., Bashford, J. A., Jr., Cooley, J. M., and Brubaker, B. S. (2001). "Detection of
9 acoustic repetition for very long stochastic patterns," Percept Psychophys **63**, 175-182.
- 10 Wright, B. A., and Zhang, Y. (2009). "Insights into human auditory processing gained from
11 perceptual learning," in *The Cognitive Neurosciences*, edited by M. S. Gazzaniga (MIT
12 Press), pp. 353-365.

TABLE

<i>Tone-cloud density</i>	<i>RT</i>			<i>RefRT</i>		
<i>(tones/s/oct.)</i>	<i>F</i>	<i>p</i>	half-life	<i>F</i>	<i>p</i>	half-life
4×2	9.21	.002	2.4	0.99	ns	—
8×4	6.26	.009	5.4	0.75	ns	—
16×8	11.46	.001	1.8	2.50	ns	—
32×16	0.39	ns	—	1.09	ns	—
2×256	1.56	ns	—	7.18	.005	18.0
128×64	1.95	ns	—	8.84	.002	1.8
512×256	0.98	ns	—	9.76	.002	0.79
matched noise	10.22	.001	7.3	1.09	ns	—

Table 1. Statistics for the time course of hit rates for RT and RefRT stimuli over trials (see also Figure 3). Where there were significant changes in hit rate over the series of trials, estimates of the half-life of the exponential fit are also listed.

FIGURE CAPTIONS

Figure 1. Schematics and cochleagrams of some example tone clouds. **A:** Each black rectangle represents a tone pip, with gray level indicating amplitude (with particularly light gray at the limits of frequency). The red lines indicate the time-frequency grid with which these tone clouds were designed, here with parameters of 2 columns per second and 1 row per octave. Exactly one tone pip starts in each cell of the grid, but because of the random offset position within its cell, its offset may stray into the subsequent cell. **B:** Cochleagram of the stimulus outlined in Panel A, with local intensity represented as brightness. **C and D:** As for panels A and B, but for a denser stimulus generated with more columns/s. **E and F:** As for panels C and D, but for a yet denser stimulus generated with more rows/octave. **G:** Cochleagram of a noise spectrally matched to the tone clouds. **H:** Cochleagram of one of the densest tone clouds used in the experiments (512 columns/s $\times$ 256 rows/octave).

Figure 2. Mean repetition-detection sensitivities observed in Experiment 1 for tone clouds with 15 different parameter sets and for the two noise stimuli. Error bars are 95% confidence intervals centered on the mean.

Figure 3. Mean repetition-detection sensitivities observed in Experiment 2, for tone clouds containing a single within-trial repetition but presented on only one trial (RT) and tone clouds containing a within-trial repetition that re-occurred during an experimental block (RefRT). Error bars are 95% confidence intervals centered on the mean.

Figure 4. Time courses of the hit-rate across blocks for tone clouds of different densities and constructions (**A–G**) and spectrally matched noise (**H**) separated into RTs (black circles) and RefRTs (red circles). The x -axis represents the presentation order within each type of stimulus across a block, not counting the intervening trials of other stimulus types. The fitted lines are exponential where the hit-rate was observed to change significantly, or straight averages otherwise (see text for details), and the shaded areas around these indicate asymptotic 95% confidence intervals centered on the mean. The hit rates on the first presentation of each stimulus type were expected to be the same, whereas any subsequent difference between RT and RefRT is indicative that re-occurrence induced learning of the RefRTs.

Figure 5. A: Modulation depths at 4 Hz at a center frequency of 400 Hz for the tone clouds and noise presented in Experiment 2, including “standard” tone clouds, whose rows and columns double proportionately, and “special” tone clouds with just two columns and many rows (see text for details). The gray data and corresponding gray axis reproduce the mean repetition-detection sensitivities of RT from Experiment 2, originally shown in Figure 3. **B:** Heat-map showing the strength of correlation between the mean RT repetition detection measured behaviorally (see panel A) and average modulation depths obtained from the model over a range of center frequencies and modulation frequencies. All positive correlations are shown, from dark to bright. **C:** As for panel B, but with correlations based on only six points, omitting the two least-dense tone-cloud conditions.