Supplementary File 1

A Dataset and Analysis Framework for Large-Scale Synthesizer Preset Generation

Joseph M. Cameron and Alan F. Blackwell

This supplementary file accompanies the main article and contains analyses and figures referenced from Sections 5 and 7 of the main text: a per-descriptor view of the principal-component–descriptor correlations (Section S1), and a comparison of factory and random presets in PCA latent space (Section S2). All analyses use the same data, principal-component decomposition, and analysis scripts as the main article; the scripts and full-resolution figures are released in the dataset repository.

S1 Per-Descriptor View of the PC–Descriptor Correlations

Section 5.3 of the main article correlates the scores of each of the first 20 principal components of the control space with the seven AudioCommons timbral descriptors (Pearce et al., 2019), and summarizes the result as a single PC–descriptor heatmap. Figure S1 restates the same information sliced by descriptor, with one panel per AudioCommons descriptor and bars ordered by absolute correlation, for readers who prefer to inspect each descriptor's PC profile directly. The three groupings reported in the main text are visible panel by panel: *brightness*, *sharpness* and *roughness* all peak on PC2; *warmth*, *depth* and *boominess* all peak on PC12; *hardness* peaks on PC4.

S2 Comparing Factory and Random Presets in Latent Space

Figure S2 compares factory and random presets in the PC1–PC2 plane of the control-space PCA described in Section 5.3 of the main article. Factory presets form tight clusters along a curved manifold, consistent with a highly constrained expert design space; random presets occupy a far broader but still non-uniform distribution shaped by the nonlinear interactions of oscillators, filters, envelopes, and modulation sources that restrict which parameter combinations yield stable, audible signals. The overlay in Figure S3 sharpens the picture: factory presets appear as concentrated islands inside the larger random cloud. Some high-density regions of random exploration contain few or no factory sounds (legal but musically rare timbres such as unbalanced envelopes, extreme filter configurations, atypical modulation routings); conversely, the most populated region of random presets overlaps a principal factory cluster, suggesting that the synthesizer itself induces certain "default" timbral tendencies that human designers and uniform search both gravitate toward. The combined dataset thus exposes both the constrained, musically validated factory manifold and a broader region of legal-but-rare timbres, supporting research on timbral interpolation, latent-space steering, and interpretable generative control.

As discussed in the main article, this distributional relationship is expected by construction: the random presets are generated by factory-conditioned mutation (Section 3.4 of the main article), so they remain anchored to the factory priors while expanding coverage of the parameter space.

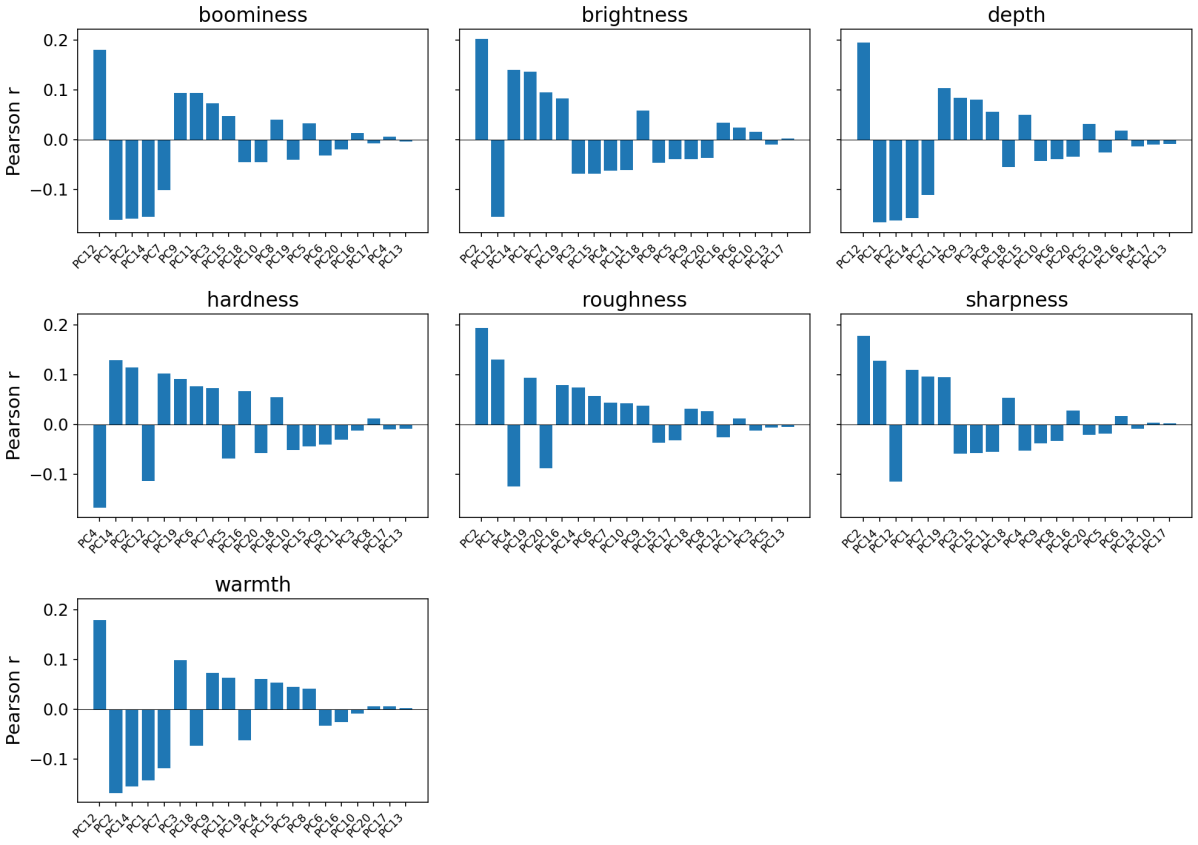


Figure S1: Per-descriptor PC correlations as a single consolidated grid: one panel per AudioCommons descriptor, with Pearson's r against the first 20 PCs, bars ordered by absolute correlation. The same information as the PC–descriptor heatmap in the main article (Section 5.3) sliced by descriptor.

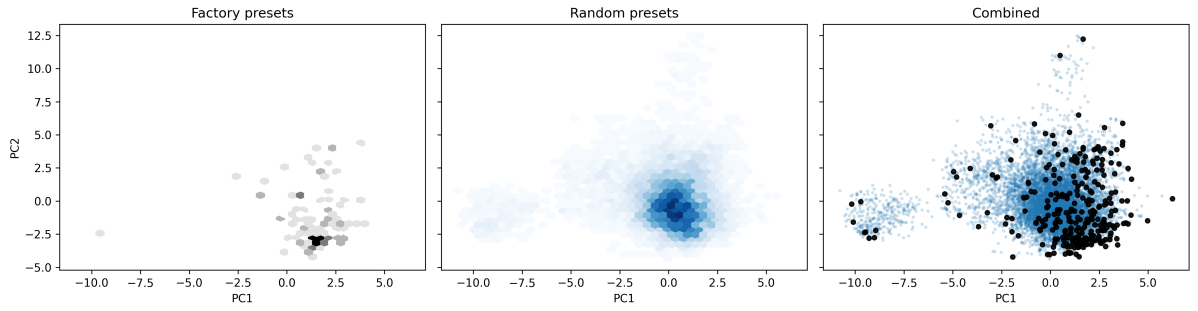


Figure S2: Comparison of factory (left) and random (center) presets in the PC1–PC2 latent space, along with a side-by-side juxtaposition (right). Factory presets form tight, structured clusters, whereas random presets occupy a broader but non-uniform region shaped by synthesizer nonlinearities.

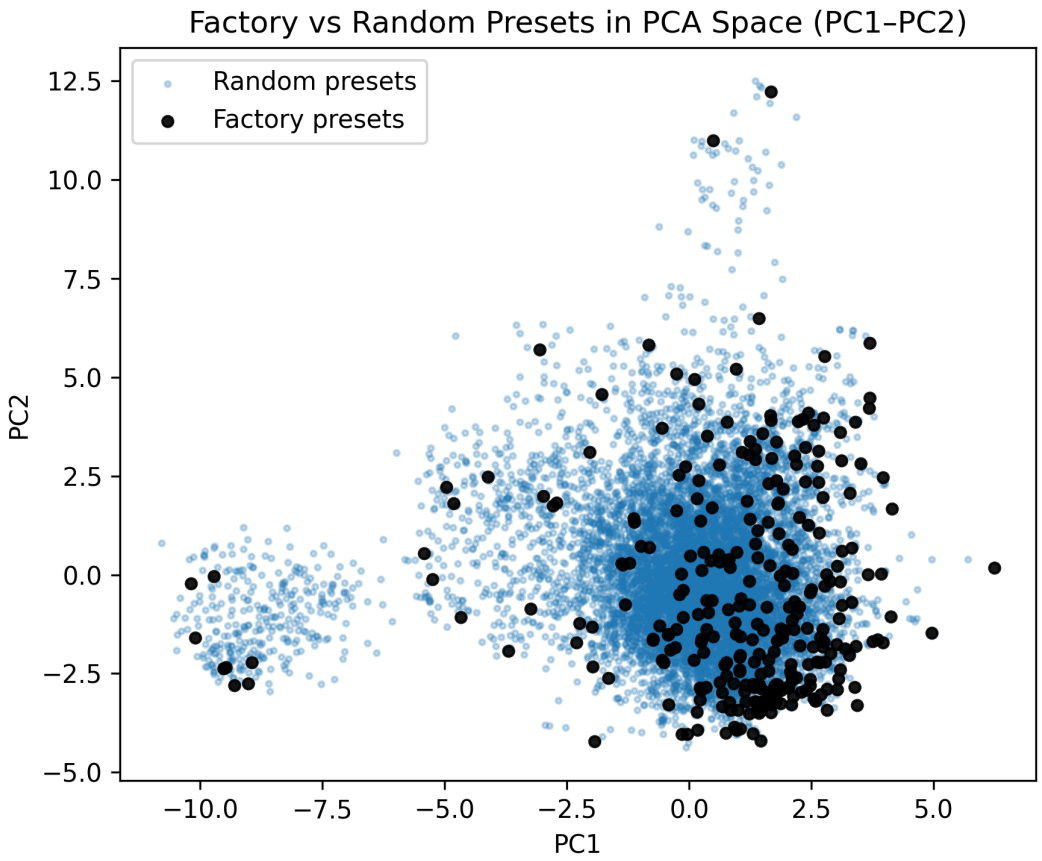


Figure S3: Overlay of factory presets (black) on the random preset cloud (blue) in PC1–PC2 space. Factory presets appear as concentrated islands embedded within the larger random distribution, revealing a human-centered design manifold embedded inside a broader parameter space.

References

Andy Pearce, Saeid Safavi, Tim Brookes, Russell Mason, Wenwu Wang, and Mark Plumbley. Deliverable d5.8: Release of timbral characterisation tools for semantically annotating non-musical content. Technical Report D5.8, Audio Commons: An Ecosystem for Creative Reuse of Audio Content, University of Surrey, 2019. Available from https://audiocommons.github.io/materials/ (accessed 2026-09-16).