EPFL

Modelling Hierarchical Key Structure With Pitch Scapes

Robert Lieck
research@robert-lieck.com

Martin Rohrmeier
martin.rohrmeier@epfl.ch

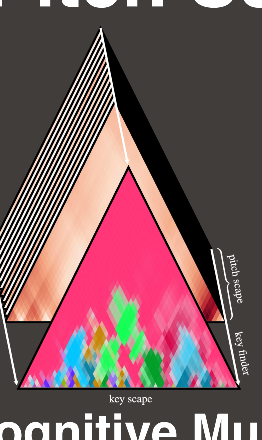

Digital and Cognitive Musicology Lab
EPFL, Switzerland

Get the code!

github.com/robert-lieck/pitchscapes

Key-Scape Plots

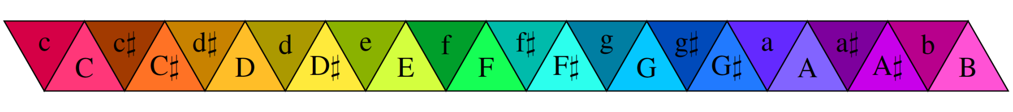

The harmonic structure is one of the most important properties of a piece in Western classical music.

Key-scape plots [1,2] visualise this hierachcical structure by mapping each subsection of a piece to a colour in a triangular plot. But they do not allow for quantitativ statistical evaluations based on large corpora of musical pieces.

We describe a novel representation and Bayesian model of the underlying statistics, which allows for unsupervised discovery of prototypical modulation plans and automatic anaylsis of the hierarchical key structure across a broad range of musical styles.

[1] Sapp CS (2001) Harmonic Visualizations of Tonal Music. In: ICMC.
[2] Sapp CS (2005) Visual hierarchical key analysis. Computers in Entertainment.

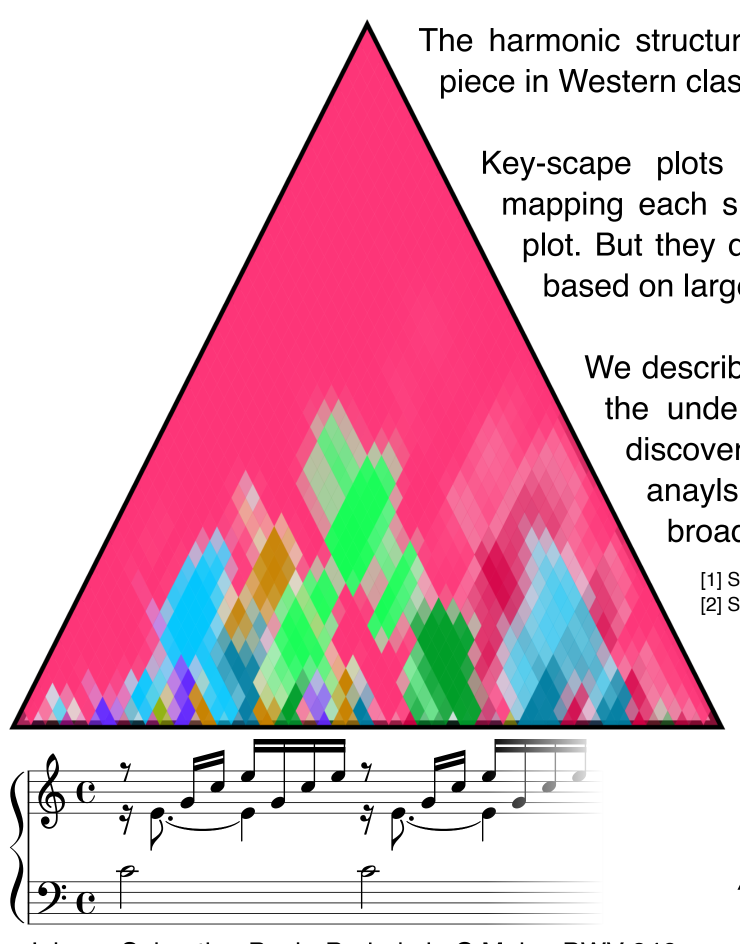

Johann Sebastian Bach, Prelude in C Major, BWV 846

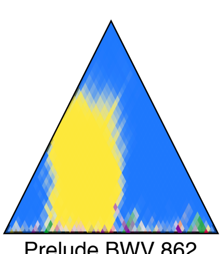

Prelude BWV 862

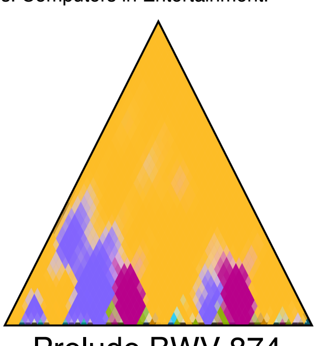

Prelude BWV 874

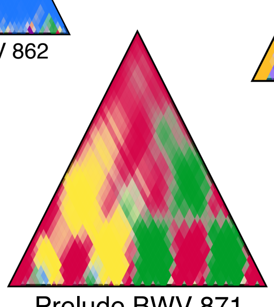

Prelude BWV 871

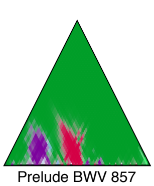

Prelude BWV 857

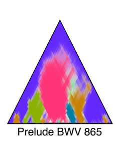

Prelude BWV 865

Pitch-Scapes

The statistics behind a key-scape plot can be modelled as a 3D continuous function.

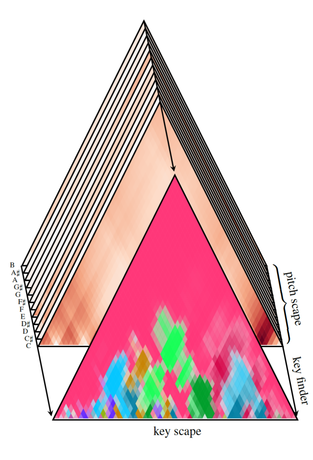

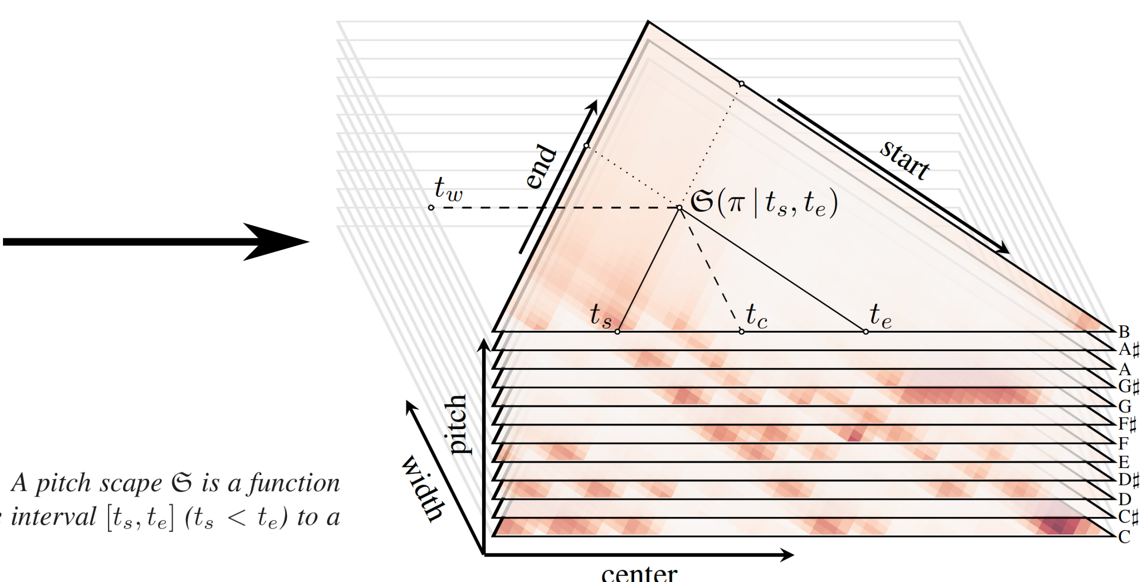

Definition 1 (Pitch Scape)**.** *A pitch scape $\mathfrak{S}$ is a function that maps each proper time interval $[t_s, t_e]$ ($t_s < t_e$) to a pitch class distribution*

$$\mathfrak{S} : \mathbb{R} \times \mathbb{R} \to [0,1]^{12} , \quad \sum_{\pi=0}^{11} \mathfrak{S}(\pi \mid t_s, t_e) = 1 . \qquad (1)$$

Bayesian Estimates

For a particular piece, we can compute a Baysian estimate of the pitch-scape by first computing the probability density over pitch classes and then integrating over the respective time interval. The prior counts *c* determine the effective time scale and ensure normalisation.

Definition 2 (Pitch Class Density)**.** *The pitch class density $\delta(\pi \mid t)$ for pitch class π at time t corresponds to the normalised pitch class counts over all tones that sound at time t*

$$\delta(\pi \mid t) := \frac{1}{\max\{1, |T_t|\}} \sum_{\tau \in T_t} [\![\tau \bmod 12 = \pi]\!] , \qquad (2)$$

where T_t is the multiset of all tones (as integers in MIDI pitch representation) sounding at time t; $[\![\cdot]\!]$ is the Iverson bracket, which equals 1 if its argument is true and 0 otherwise; and the max *avoids division by zero for silent parts where $T_t = \emptyset$ is the empty set.*

Definition 3 (Pitch Scape Estimate)**.** *The posterior estimate of the pitch scape $\mathfrak{S}(\pi \mid t_s, t_e)$ for pitch class π and time interval $[t_s, t_e]$ is*

$$\mathfrak{S}(\pi \mid t_s, t_e) := \underbrace{\frac{1}{t_e - t_s + 12c}}_{\text{normalisation}} \Big[\overbrace{c}^{\text{prior counts}} + \underbrace{\int_{t_s}^{t_e} \delta(\pi \mid t)\, dt}_{\text{overall pitch class counts}} \Big] , \qquad (3)$$

where the integral over the pitch class density computes the overall pitch class counts, $c \geq 0$ specifies the prior counts, and the leading term ensures proper normalisation.

Prototypes

A probabilistic prototype for pitch-scapes is defined as a point-wise Dirichlet distribution, with the function $\alpha > 0$ representing the parameters at each point. This allows computing the log-likelihood for any piece given a particular prototype α.

Definition 4 (Prototype)**.** *Given a function*

$$\alpha : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+^{12} \qquad (4)$$

that maps each proper time interval $[t_s, t_e]$ ($t_s < t_e$) to a vector with positive entries, a prototype is defined as the point-wise Dirichlet distribution with parameter vector α. The likelihood of observing a pitch class distribution Π for the interval $[t_s, t_e]$ given α is

$$p(\Pi \mid \alpha, t_s, t_e) = \mathrm{Dir}(\Pi; \alpha(t_s, t_e)) . \qquad (5)$$

The log-likelihood of observing a full pitch scape $\mathfrak{S}$ given α is

$$\log p(\mathfrak{S} \mid \alpha) = \frac{2}{T^2} \iint_{0 \leq t_s < t_e \leq T} \log \mathrm{Dir}\big(\mathfrak{S}(t_s, t_e); \alpha(t_s, t_e)\big) dt_s dt_e , \qquad (6)$$

where T is the duration of the piece.

Fourier Representation

A prototype α is represented via a finite Fourier series with parameters θ

$$\alpha_\pi^{(\theta,\tau)}(t_s, t_e) := e^{\widetilde{\alpha}^{(\theta,\tau)}(\bar{t}_c, \bar{t}_w, \pi)} \qquad (8)$$

$$\widetilde{\alpha}^{(\theta,\tau)}(\mathbf{x}) = \sum_{\mathbf{n}} \theta_{\mathbf{n}} e^{2\pi i \mathbf{k}_{\mathbf{n}} \cdot \mathbf{x}} , \qquad (10)$$

$$\mathbf{x} := (\bar{t}_c, \bar{t}_w, \pi) \qquad n_c \in \{-N_c, \dots, N_c\} \qquad (11)$$

$$\mathbf{n} := (n_c, n_w, n_\pi) \qquad n_w \in \{-N_w, \dots, N_w\} \qquad (12)$$

$$\mathbf{k}_{\mathbf{n}} := \left(\sigma_c n_c, \sigma_w n_w, \tfrac{n_\pi + \tau}{12}\right) \qquad n_\pi \in \{-6, \dots, 6\} . \qquad (13)$$

Mixture Model

Multiple prototypes are combined in a mixture model to describe an entire corpus of musical pieces. Each piece is assumed to be generated by a specific prototype (or cluster) *c* with transposition *τ*.

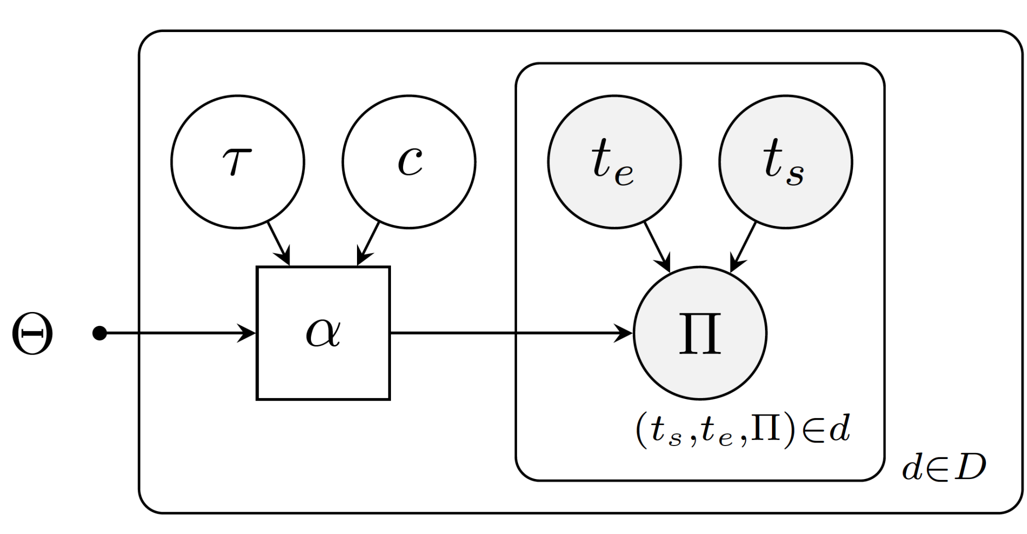

Hierarchical Clustering

Using the mixture model, we performed unsupervised hierarchical clustering on a corpus of 155 Baroque pieces by J.S. Bach, G.F. Händel, and D. Scarlatti. The final clusters empirically confirm common prototypes postulated in music theory.

Visualisations of the prototypes are realised with conventional key-scape plots. Note that this only displays a small part of the nuanced information contained in the underlying pitch scapes. For prototypes, the indicated keys are only valid up to common transposition, as the prototypes themselves are transposition invariant.

circle of fifths

Colouring along the circle of fifths maps related keys to similar colours.

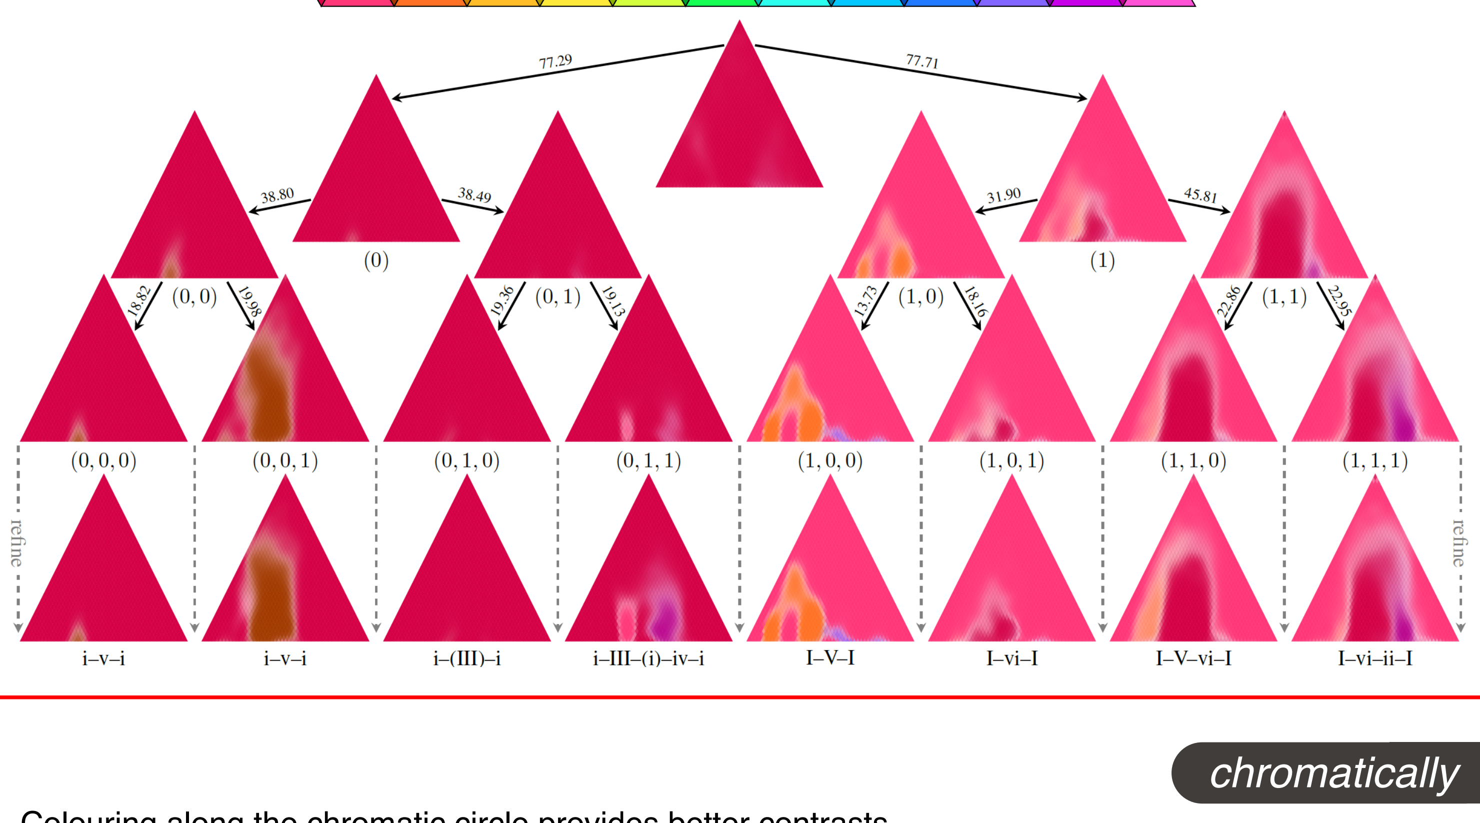

chromatically

Colouring along the chromatic circle provides better contrasts.

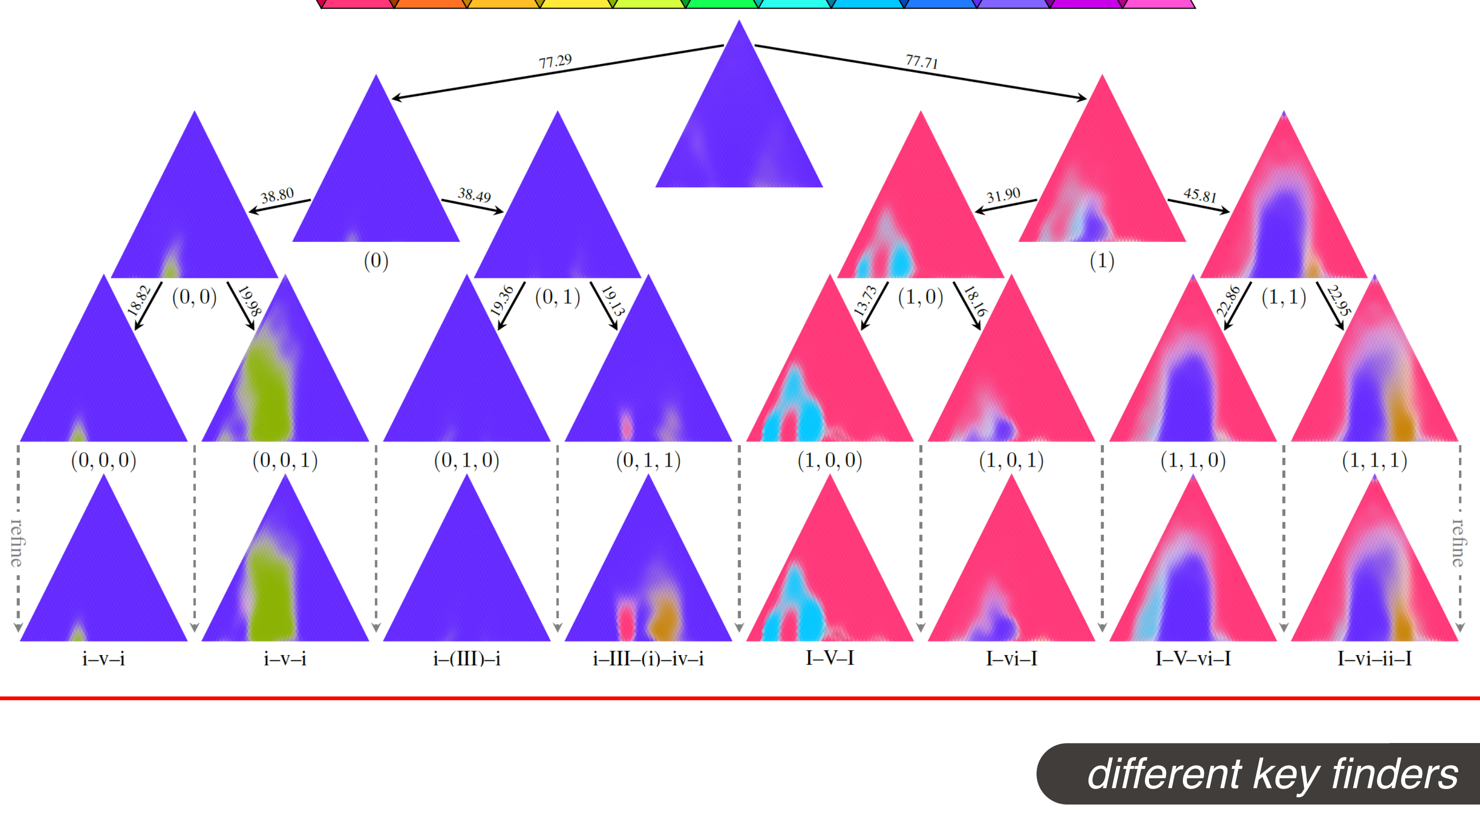

different key finders

Different key-finding algorithms result in different key estimates for the same underlying pitch scapes, which allows to reveal some more information.

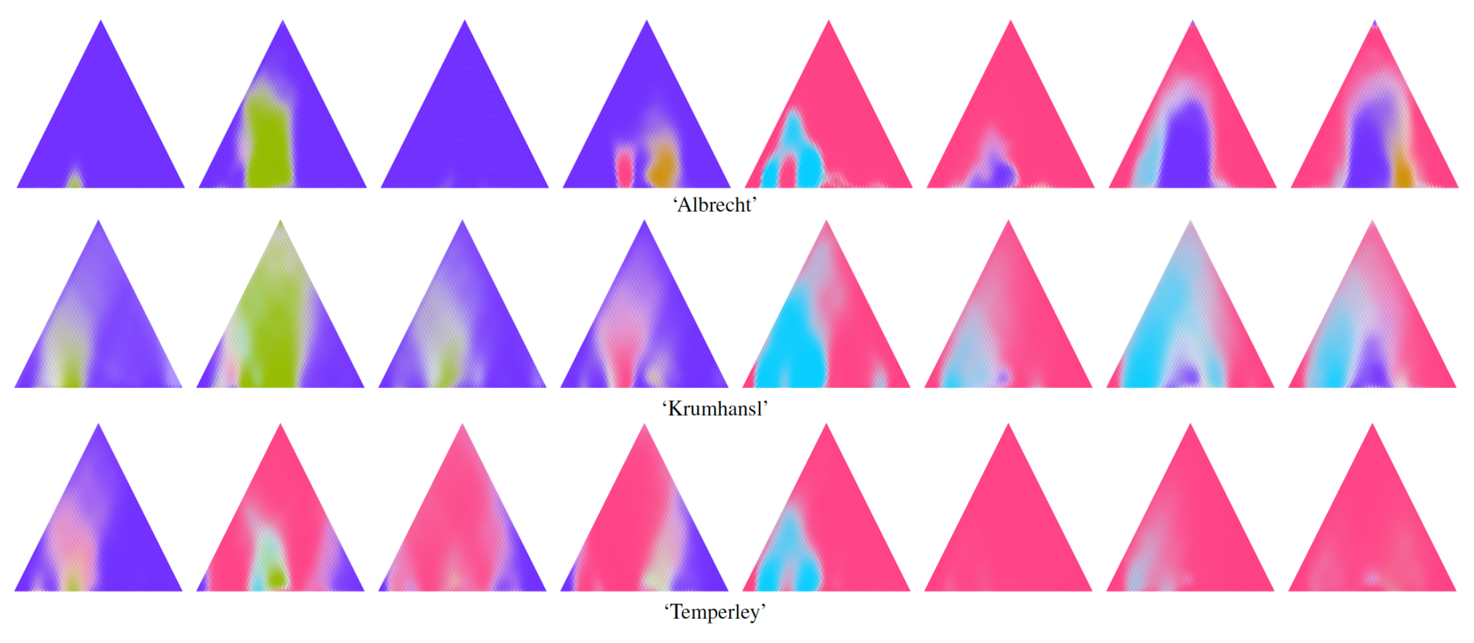

Acknowledgments

European Research Council
Established by the European Commission

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 760081 – PMSB.