

Visualizing Directional Dependence on Cylinders and Tori

Aurélien Nicosia

Université Laval

Abstract

Directional data arise in many scientific settings where observations are measured on a repeating angular scale. Existing circular plots and joint density displays are useful, but they do not by themselves organize conditional estimands for cylindrical and toroidal sample spaces. This paper proposes a conditional graphical grammar for visualizing directional dependence when an angular variable is paired with either a linear or another angular response. The main displays are circular topography, toroidal topography and conditional ridges. Each is tied to an explicit estimand, such as a conditional density or modal relation. Simulation scenarios and a buoy wind-wave example show how these displays reveal multimodality, seam-crossing structure and dependence that can be obscured by marginal or unfolded views. The methods are implemented in the R package **ggcircular**, which returns standard **ggplot2** objects.

Keywords: Directional data, statistical visualization, conditional density, toroidal data, R.

1. Introduction

Directional observations require geometric care because their support is periodic. Values near 0 and 2π are adjacent, not opposite endpoints of a linear interval. Directional statistics therefore treat angular observations on S^1 rather than on $\mathbb{R}$ (Mardia and Jupp 2000; Jammalamadaka and SenGupta 2001; Pewsey and García-Portugués 2021).

When an angular variable is observed with a real response, the sample space is the cylinder $S^1 \times \mathbb{R}$. When it is observed with another angle, the sample space is the torus $S^1 \times S^1$.

Classical circular graphics remain useful. Rose diagrams, circular histograms, circular densities, polar scatterplots and unfolded angle plots show marginal concentrations and simple trends. They are less direct for questions such as: how does the full distribution of X change with Θ ; does Φ follow a conditional ridge as Θ changes; is a conditional mean hiding two modes; does a structure cross the seam $0/2\pi$; and can a dependence be visible even when the marginal angular distributions are nearly uniform?

These questions are conditional rather than purely joint or marginal. A marginal display describes $f_\Theta(\theta)$ or $f_X(x)$ separately. A joint display describes where observations concentrate in (θ, x) or (θ, ϕ) , but this concentration can be driven by the marginal frequency of the conditioning angle. The estimand used throughout this paper is instead the conditional distribution, such as $f(x \mid \theta)$ or $f(\phi \mid \theta)$, so that angular slices are compared on a common conditional scale.

This article proposes a graphical grammar for these conditional questions. The term grammar is used in the sense of [Wilkinson \(2005\)](#) and [Wickham \(2010\)](#); a graph is built from data, mappings, statistical transformations, geometries, scales and themes. The contribution here is to make the statistical transformation conditional and geometry-aware for cylinders and tori.

2. Contribution and Positioning

This paper establishes a conditional graphical grammar for visualizing directional dependence. Rather than relying on marginal or joint displays alone, it formalizes statistical graphics as direct representations of conditional random variables on cylinders and tori.

The proposed grammar has four linked contributions. First, it gives a common probabilistic formalization for $S^1 \times \mathbb{R}$ and $S^1 \times S^1$ data. Second, it connects graphical displays to explicit estimands, especially conditional densities, conditional moments and modal ridges. Third, it uses simulation scenarios to document descriptive visual signatures under independence and structured alternatives. Fourth, it provides a reproducible R implementation in **ggcircular**, compatible with **ggplot2**.

This framing separates the proposed work from existing ingredients. Circular visualizations, including area-proportional variants, are well established ([Xu and Wang 2020](#)), and the R package **circular** provides a broad software baseline for circular statistics ([Agostinelli and Lund 2025](#)). Kernel smoothing for circular predictors and circular responses has been studied by [Di Marzio et al. \(2009\)](#) and [Di Marzio et al. \(2013\)](#), while directional-linear density estimation has been studied by [García-Portugués et al. \(2013\)](#). Toroidal density estimation and ridge methods provide further anchors ([Di Marzio et al. 2011](#); [García-Portugués and Prieto-Tirado 2023](#)). Toroidal methods also exist for clustering and dimension reduction, including **ClusTorus** ([Hong and Jung 2022](#)). Software such as **ggcircular** provides a **ggplot2** extension for marginal circular and directional graphics ([Nicosia 2026](#)).

These works provide the statistical and graphical ingredients on which the present paper builds. The added contribution is to assemble them into a single conditional workflow that specifies estimands, expected null signatures, visual displays and **ggplot2** layers for both $S^1 \times \mathbb{R}$ and $S^1 \times S^1$. The graphical comparison is given after the displays are introduced, so that the reader can compare existing views and the proposed views on concrete examples.

3. Statistical Framework

Throughout the paper, angles are represented on $[-\pi, \pi)$ for plotting but all smoothing, distances and integrations are computed modulo 2π . The signed circular difference is

$$d_{S^1}(a, b) = \text{atan2}\{\sin(a - b), \cos(a - b)\},$$

and angular errors use $|d_{S^1}(a, b)|$. The seam is therefore a display choice, not a boundary in the statistical sample space.

3.1. Cylindrical data

For circular-linear data, observe (θ_i, x_i) , $i = 1, \dots, n$, with $\theta_i \in S^1$ and $x_i \in \mathbb{R}$. A product-kernel density estimator can be written as

$$\hat{f}(\theta, x) = \frac{1}{n} \sum_{i=1}^n K_\kappa(\theta - \theta_i) L_h(x - x_i),$$

where K_κ is a normalized circular kernel on S^1 and L_h is a normalized linear kernel on $\mathbb{R}$. The implementation uses scaled von Mises weights in the angular coordinate and a Gaussian kernel in the linear coordinate; row normalization is then applied on the display grid. The angular coordinate has no boundary in the statistical calculation, but local smoothing bias and finite-grid truncation in the response coordinate can still matter for small samples or rapidly changing conditional shapes. The main object for visualization is the conditional density

$$\hat{f}(x \mid \theta) = \frac{\hat{f}(\theta, x)}{\int_{\mathbb{R}} \hat{f}(\theta, u) \, du}.$$

When L_h integrates to one, the same estimator can be written as a weighted mixture in the response coordinate:

$$\hat{f}(x \mid \theta) = \sum_{i=1}^n w_i(\theta) L_h(x - x_i), \quad w_i(\theta) = \frac{K_\kappa(\theta - \theta_i)}{\sum_{j=1}^n K_\kappa(\theta - \theta_j)}.$$

This expression makes the conditional normalization explicit: angular slices are compared after removing the marginal density of the conditioning angle. Moment summaries for companion displays are given in the appendix; the main cylindrical display uses the full conditional density.

3.2. Toroidal data

For circular-circular data, observe (θ_i, ϕ_i) with $\theta_i, \phi_i \in S^1$. A density estimator on the torus is

$$\hat{f}(\theta, \phi) = \frac{1}{n} \sum_{i=1}^n K_{\kappa_\theta}(\theta - \theta_i) K_{\kappa_\phi}(\phi - \phi_i).$$

The notation κ_θ denotes smoothing in the conditioning angle and κ_ϕ for smoothing in the response angle. Both circular kernels are normalized on S^1 in the statistical definition; the implementation uses stable scaled von Mises weights and normalizes the conditional density on the grid. Conditional likelihood scores include the response-angle normalizing constant so that candidate κ_ϕ values are comparable.

The conditional density is

$$\hat{f}(\phi \mid \theta) = \frac{\hat{f}(\theta, \phi)}{\int_0^{2\pi} \hat{f}(\theta, u) du}.$$

Equivalently,

$$\hat{f}(\phi \mid \theta) = \sum_{i=1}^n w_i(\theta) K_{\kappa_\phi}(\phi - \phi_i), \quad w_i(\theta) = \frac{K_{\kappa_\theta}(\theta - \theta_i)}{\sum_{j=1}^n K_{\kappa_\theta}(\theta - \theta_j)}.$$

The local circular moment is estimated by

$$\widehat{m}(\theta) = \frac{\sum_i K_{\kappa}(\theta - \theta_i) e^{i\phi_i}}{\sum_i K_{\kappa}(\theta - \theta_i)}.$$

The direction and concentration summaries are

$$\hat{\mu}_\Phi(\theta) = \text{Arg}\{\widehat{m}(\theta)\}, \quad \hat{\rho}_\Phi(\theta) = |\widehat{m}(\theta)|.$$

A high value of $\hat{\rho}_\Phi(\theta)$ indicates local angular concentration. A low value can indicate dispersion, symmetry or conditional multimodality.

The modal ridge uses the same conditional density; its formal definition is given with the conditional ridge display below. Together these estimands define the toroidal topography and conditional ridge used in the main figures.

3.3. Marginal support

Conditional displays should not be interpreted strongly where the conditioning angle is weakly represented. For a grid point θ_g , the implementation computes

$$S(\theta_g) = \sum_{i=1}^n K_{\kappa_s}(\theta_g - \theta_i), \quad R(\theta_g) = \frac{S(\theta_g)}{\max_{g'} S(\theta_{g'})}.$$

A point is flagged when $R(\theta_g) < \tau$, where the software default is $\tau = 0.15$ and the supplementary example uses $\tau = 0.25$. The threshold is a relative effective-support rule rather than an asymptotic confidence rule: for fixed smoothing, the variance of a local kernel estimate increases as the available kernel mass decreases. Thus τ should be interpreted as a user-controlled tolerance for unstable local estimates and checked

by sensitivity analysis in sparse telemetry settings. This flag is a diagnostic warning, not a formal test. It indicates that a conditional estimate is based on little angular information.

A practical interpretation follows from the usual inverse-kernel-mass variance relationship for kernel estimators (Wand and Jones 1995). Holding the response kernel and local conditional variability fixed, the standard error at θ_g is approximately inflated by $1/\sqrt{R(\theta_g)}$ relative to the best-supported direction. The default $\tau = 0.15$ therefore permits an approximate inflation of 2.58, whereas the more conservative $\tau = 0.25$ flags regions once the inflation reaches about 2. A value of 0.15 is recommended for an exploratory first view and 0.25 when a more conservative support screen is appropriate. If a substantive reading changes between these values, both results should be reported as a sensitivity analysis. Intermediate values are also legitimate: the software vignette uses $\tau = 0.20$ as an illustrative compromise between these two reference points, not as an additional recommended cutoff.

3.4. Smoothing selection

The concentration and bandwidth parameters can be supplied by the user or chosen from a candidate grid by cross-validated conditional likelihood. The first row of the ranked selection table is used as the likelihood-selected smoothing choice, but the article reports sensitivity because conditional visual readings can change when the data are sparse or multimodal. In mixtures, the selected smoothing can merge nearby branches or split noisy local features depending on the candidate grid and sample size; sensitivity plots are therefore part of the intended workflow. The cross-validated score optimizes predictive conditional density within the supplied smoothing grid; it does not identify the minimal mixture order and can be sensitive when finite-state or mixture approximations with different orders fit similar conditional shapes. Full formulas and software details are given in the appendix.

4. Graphical Grammar

The grammar contains six implemented displays, but the main article focuses on three displays that carry the statistical argument: circular topography, toroidal topography and conditional ridge. Phase loom, toroidal flow and statistical orbit are companion summaries for binned masses or near-unimodal conditional moments; they are documented in the appendix and in the software. Throughout the article, white lines drawn over a topography are equally spaced isocontours of the conditional density encoded by the colour field; they are not observed trajectories.

4.1. Circular topography

The circular topography represents $\hat{f}(x | \theta)$ directly. It is useful when the full conditional shape is more informative than a mean curve. In Figure 1, the high-density band follows a smooth periodic trajectory. Under independence, the rows of the topography should have similar shapes across θ .




Figure 1: Circular Topography for the same smooth cylindrical scenario. The high-density band follows a smooth function of θ . The white lines are equally spaced isocontours of the same estimated conditional density represented by the colour field; they are not observed trajectories.

4.2. Toroidal topography

The toroidal topography unfolds the torus into a periodic rectangle and represents $\hat{f}(\phi | \theta)$. In Figure 2, the conditional density follows a wrapped relation of the form $\Phi \approx 2\Theta \pmod{2\pi}$.




Figure 2: Toroidal Topography for a doubling scenario. The conditional density forms wrapped diagonal bands. The white lines are equally spaced isocontours of the same estimated conditional density represented by the colour field; they are not observed trajectories.

4.3. Conditional ridge

The conditional ridge follows

$$\hat{r}(\theta) = \operatorname{argmax}_{\phi \in S^1} \hat{f}(\phi \mid \theta).$$

The curve is coloured by $\hat{\rho}_{\Phi}(\theta)$, the local resultant length of the responses around the conditioning angle. This colour is a resultant concentration summary, not a branch-specific mode height. It can be low when the conditional density is symmetric or multimodal even if each mode is sharp. In Figure 3, the ridge follows a diagonal relation and the nearly uniform warm colour indicates high local concentration across the displayed domain.









Figure 3: Conditional Ridge for a diagonal scenario. The curve follows the modal ridge $\hat{r}(\theta)$; the nearly uniform warm colour indicates high local concentration $\hat{\rho}_{\Phi}(\theta)$ across the displayed domain.

The implemented ridge extraction algorithm is deliberately simple. It evaluates $\hat{f}(\phi \mid \theta_g)$ on a regular grid, selects the grid value with maximum conditional density for each θ_g , breaks the displayed line when adjacent selected values jump by more than $\pi/2$, and reports the local resultant length at the same θ_g . An exact tie means that two or more grid values attain the computed global maximum; it is resolved by selecting the first such grid value. This rule is reproducible and useful for a dominant ridge, but the selected branch is not stable when distinct modes are exactly or nearly tied. The implementation therefore identifies distinct circular local maxima and reports their number, the height ratio of the second-highest to the highest local maximum, and an ambiguity flag. A near tie is defined through this ratio. With the default `tie_tolerance = 0.01`, a grid location is flagged when the ratio is at least 0.99. The diagnostic does not change the selected first maximum and does not attempt to track multiple branches.

Ridges must be interpreted carefully. They summarize a local mode, not the entire conditional distribution. In multimodal regions, a single ridge can mask ambiguity. Near seams, a ridge may appear discontinuous because of the chosen parametrization

rather than because the underlying structure is discontinuous. A conditional ridge should not be presented as an isolated diagnostic when conditional dispersion is high, local concentration is low or multimodality is plausible or the near-tie diagnostic is flagged. In those cases it should be shown with the toroidal topography background or with an equivalent display of conditional mass, never as an isolated ridge.

Table 1: Main visual positioning of the proposed conditional displays. Detailed companion displays are reported in the appendix.

View	Existing display	Proposed display	Main limitation	Proposed advantage
Circular topography			Joint relief mixes conditional shape with the marginal density of Θ .	Row normalization targets $\hat{f}(x \theta)$ and makes angular slices comparable.
Toroidal topography			Joint density does not organize the reading by conditioning angle.	The surface directly targets $\hat{f}(\phi \theta)$.
Conditional ridge			A ridge-only curve can hide surrounding mass and branch ambiguity.	The ridge is read with conditional density, concentration and near-tie diagnostics.

5. Simulation Study

The simulation study is a descriptive controlled check. It asks whether the proposed displays show conditional structure that classical unfolded or marginal displays can hide. It is not intended to establish a formal testing procedure or to compare statistical power across displays.

Under cylindrical independence,

$$f(\theta, x) = f_{\Theta}(\theta)f_X(x), \quad f(x | \theta) = f_X(x).$$

The circular topography should repeat the same profile of X for all angles, the statistical orbit should show approximately constant mean and spread, and the phase loom should send each angular sector to similar classes of X .

Under toroidal independence,

$$f(\phi | \theta) = f_{\Phi}(\phi).$$




Figure 4: Cylindrical comparison. Left: unfolded scatterplot with a single smooth trend. Right: estimated conditional density $\hat{f}(x | \theta)$ revealing two conditional modes.

The toroidal topography should show a horizontal band or a relief that does not vary systematically with θ , and a conditional ridge should follow a marginal mode of Φ rather than a function of θ .

The protocol contains five cylindrical scenarios and five toroidal scenarios, crossed with sample size and smoothing grids. The full data-generating mechanisms, diagnostic definitions and smoothing sensitivity summaries are given in the appendix. In every scenario, including the independence and structured alternatives, the conditioning angle is generated as $\Theta \sim U(0, 2\pi)$.

Display-specific diagnostics are compared with matched independence baselines, in the spirit of visual diagnostic calibration for exploratory graphics (Buja et al. 2009). For each space, display, sample size and smoothing value, an independent null simulation estimates the 95th percentile of the diagnostic statistic. Evaluation replicates are then compared with this held-out threshold. These diagnostics summarize visual behavior in controlled scenarios; they are not formal tests for scientific inference.

The comparisons in Figures 4 and 5 illustrate why conditional displays are needed. In the cylindrical case, a single smooth trend compresses two conditional branches. In the toroidal case, a circular mean can fall between two branches, whereas the conditional density preserves the bifurcation.

The joint and conditional views differ most clearly when the conditioning angle has a non-uniform marginal distribution. A joint KDE remains useful for describing where observations concentrate in (θ, x) , but its relief is partly driven by the marginal density of Θ . Conditional topography renormalizes each angular slice and targets the changing shape of $X | \Theta = \theta$ more directly. The full joint-KDE comparison is shown in the appendix.




Figure 5: Toroidal comparison. Left: unfolded scatterplot and circular mean trend. Right: estimated conditional density $\hat{f}(\phi | \theta)$ preserving two conditional branches.

Table 2: Thresholded descriptive detection rates by display. The false-positive column uses matched independence scenarios; the alternative column pools structured scenarios. Counts are detections over evaluation trials.

Space	Display	False-positive rate	Alternative detection
Cylindrical	Circular Topography	0.067 (107/1600)	1.000 (11196/11200)
Cylindrical	Phase Loom	0.071 (113/1600)	0.987 (11057/11200)
Cylindrical	Statistical Orbit	0.062 (100/1600)	1.000 (11200/11200)
Toroidal	Conditional Ridge	0.049 (78/1600)	1.000 (13600/13600)
Toroidal	Toroidal Flow	0.039 (62/1600)	1.000 (13600/13600)
Toroidal	Toroidal Topography	0.043 (68/1600)	1.000 (13600/13600)

5.1. Quantitative results

The detection summary below reports thresholded descriptive detection rates. Null thresholds are estimated from 250 independent calibration replicates per matched null cell. Under independence, the empirical false-positive rate is 0.055 with exact binomial interval $[0.0505, 0.0597]$. For the cylindrical displays, the pooled false-positive rates range from 0.062 to 0.071, slightly above the nominal 0.05 level in some cells. This is expected for a finite Monte Carlo threshold and a visual diagnostic summary; the 95th percentile rule is not a calibrated hypothesis test.

Across structured alternatives, the overall detection rate is 0.998 $[0.9977, 0.9983]$. These alternatives are intentionally strong and are used to verify that the expected visual signatures appear under controlled conditions, not to rank the displays by subtle-signal power. The weakest cell is phase loom in the heteroscedastic scenario at $n = 250$ and $\kappa = 36$, with detection rate 0.812.

For toroidal scenarios, ridge recovery is evaluated with circular distances, so crossings of the $0/2\pi$ seam do not inflate the error. The mean angular RMSE of $\hat{r}(\theta)$ over the toroidal protocol is 0.103 radians. The largest cell mean is 0.254 radians in the doubling

scenario at $n = 250$ and $\kappa = 8$. This nearest-branch RMSE can understate instability in bifurcating or antipodal scenarios, so detailed ridge and branch-switching diagnostics are kept in the appendix.

6. Empirical Illustration

The empirical example uses historical meteorological observations from NDBC station 44013 for 2024. Station 44013 corresponds to the Boston Approach Lighted Buoy BF NOAA 44013 ([National Data Buoy Center 2026b](#)). The station and year were retained because the standard meteorological file provides a large number of valid wind and wave direction records over a complete recent annual cycle ([National Data Buoy Center 2024](#)). The example is intended as an empirical illustration of the graphics, not as a full meteorological model.

The standard meteorological file contains wind direction **WDIR**, wind speed **WSPD** and mean wave direction **MWD**. According to the National Data Buoy Center, **WDIR** is a true direction in degrees indicating where the wind comes from, **WSPD** is wind speed in m/s, and **MWD** is the direction from which waves associated with the dominant period come ([National Data Buoy Center 2026a](#)).

The left panel of Figure 6 is cylindrical, not toroidal: it crosses wind direction with wind speed, so the space is $S^1 \times \mathbb{R}$. The right panel is toroidal: it crosses wind direction with mean wave direction, so the space is $S^1 \times S^1$. The NDBC directions are true degrees, measured clockwise from north. They are converted to mathematical radians using

$$\theta = \frac{((90 - d) \bmod 360)\pi}{180}.$$

The raw missing-value codes are converted to **NA** before analysis. The file contains 52,665 raw records. The final displays use all valid pairs: 52,640 wind-direction and wind-speed pairs for the cylindrical panel and 17,106 wind-direction and mean-wave-direction pairs for the toroidal panel. Valid **MWD** is available for 17,108 records. The toroidal panel uses two fewer records because the preprocessing rule keeps observations with valid **WDIR**, **WSPD** and **MWD**, preserving a common validity rule across the wind-speed and wind-wave summaries. The toroidal display should therefore be read as a conditional description of the subset with valid wave direction and complete wind records.

Smoothing parameters are selected by five-fold cross-validated conditional likelihood on deterministic tuning samples of at most 5,000 valid pairs, and the selected parameters are then applied to the full valid samples. The tuning rows are taken at equally spaced positions in the ordered complete-record data. Fold labels are balanced and randomly permuted using seeds 20260615 and 20260616 for the cylindrical and toroidal panels, respectively. For the cylindrical panel, cross-validation used candidate values $\kappa \in \{8, 14, 24, 36\}$ and $h \in h_0\{0.6, 0.8, 1.0, 1.3, 1.7\}$, where h_0 is the reported rule-of-thumb bandwidth calculated on the tuning sample, and selected $(\kappa, h) = (24, 0.578)$. For the toroidal panel, cross-validation used $(\kappa_\theta, \kappa_\phi) \in \{8, 14, 24, 36\}^2$ and selected $(\kappa_\theta, \kappa_\phi) = (36, 36)$. This toroidal choice is on the upper boundary of the candidate grid, so the panel indicates a preference for relatively concentrated smoothing within




Figure 6: Empirical example from valid observations at NDBC station 44013 in 2024. Left: conditional density of wind speed given wind direction. Right: conditional density of mean wave direction given wind direction. Smoothing parameters are selected by conditional likelihood cross-validation and applied to all valid pairs.

the grid rather than a definitive global optimum. Axis labels are shown in compass order W, S, E, N, W because NDBC true degrees are converted to mathematical radians for plotting; physical interpretation remains in terms of meteorological “from” directions.

As a descriptive check against an exchangeable independence baseline, Cramér’s V was computed after discretization and compared with response permutations that preserve the observed marginals. With 999 exchangeable permutations, the wind-speed panel has statistic 0.102 and permutation $p \leq 0.001$; the wave-direction panel has statistic 0.215 and permutation $p \leq 0.001$. The same discretization with daily block permutations gives $p \leq 0.001$ for wind speed and $p \leq 0.001$ for wave direction. These values use the add-one permutation convention $(b+1)/(B+1)$; with 999 permutations, 0.001 is the smallest attainable p -value, and $p \leq 0.001$ means that no permuted statistic was at least as large as the observed statistic. Because the records are temporally ordered, these permutation results are used as diagnostics rather than as full time-series independence tests. Sensitivity to 12, 24 and 36 angular bins is reported in the reproducibility appendix.

The empirical displays show that, at this station and year, the conditional distribution of wind speed changes with wind direction and that mean wave direction is not conditionally uniform given wind direction. These patterns may reflect seasonal weather regimes, coastal geometry, fetch, storm events and measurement availability. A full environmental analysis would require temporal covariates and physical modeling.

7. Software Implementation

The methods are implemented in the R package **ggcircular**. The software role in this article is reproducibility: the figures, simulation summaries and empirical illustration are generated from documented functions rather than from ad hoc plotting code. The

main plotting functions return **ggplot2** objects, and the core displays are also available as statistical layers.

The core manuscript functions are the circular topography, toroidal topography and toroidal ridge plotting functions, with matching **ggplot2** statistical layers. The topography functions take the observed angles and responses, smoothing parameters and grid sizes, and return **ggplot2** objects or **ggplot2** layers. Companion functions implement phase looms, toroidal flows, statistical orbits, smoothing selection and diagnostic summaries.

The article figures use fixed grids to make the rendered displays reproducible: 150 by 130 for circular topography, 150 by 150 for toroidal topography and toroidal ridge backgrounds, 130 by 120 for the empirical circular-linear density, and 120 by 120 for the empirical toroidal density. Quantitative simulation metrics use coarser grids to keep the Monte Carlo protocol reproducible and computationally feasible.

The implementation also includes diagnostic helpers for simulation scenarios, bootstrap intervals, smoothing sensitivity, marginal support flags and the companion displays described in the appendix. The package includes roxygen documentation, unit tests for row normalization, seams, ridge recovery, non-finite inputs and smoothing reproducibility, plus vignettes covering cylindrical, toroidal and smoothing-support workflows.

8. Discussion

The main contribution is a change in graphical organization. The proposed displays do not introduce new directional kernels or new toroidal density estimators. They reorganize existing statistical ingredients around conditional estimands and make those estimands available as reproducible graphics. This is important because marginal and joint views can answer a different question from the one an analyst often asks: how does the distribution of one response change as an angle varies?

The approach is exploratory. Smoothing choices affect the displayed structure, low marginal support can create unstable conditional estimates, and a single ridge can hide branch ambiguity in multimodal toroidal data. Cross-validated smoothing selects a predictive display within the specified candidate grid; it does not guarantee recovery of a minimum number of mixture components. In larger analysis pipelines, including emission-distribution diagnostics for hidden Markov models of animal movement, the displays should therefore be treated as conditional graphical diagnostics of the extracted variables rather than as validation of the upstream model. The simulation diagnostics support interpretation of the graphics, not formal scientific testing. The bootstrap intervals used here are descriptive and can under-cover the generating truth when smoothing bias dominates sampling variation. In the NDBC example, the displays identify directional patterns worth modeling, but they do not replace a time-series or physical oceanographic analysis.

Future work should focus on adaptive smoothing, explicit multi-branch ridge extraction, temporally dependent directional data and reader studies for graphical perception. These extensions would strengthen inference and display design without changing the central point of the present article: conditional structure on cylinders and tori should be visualized through conditional estimands.

9. Reproducibility and Software Availability

The methods are implemented in **ggcircular** version 0.2.0. The exact source used for this article, the package source tarball and the public documentation are available from the tagged GitHub release at <https://github.com/AurelienNicosiaULaval/ggcircular/releases/tag/v0.2.0>. The release also provides a reproducibility supplement containing the fixed-seed analysis scripts, configuration, input data, per-replication results, generated tables and figures, and session information used for the article.

The supplement supports a fast manuscript render from the included outputs and a full recomputation of the simulation and empirical analyses. The entry point for the complete workflow is `Rscript scripts/run_full_analysis.R`; exact commands and computational requirements are documented in the accompanying reproduction README.

Computational Details

The manuscript was rendered with R version 4.5.0 (2025-04-11) using **ggcircular** version 0.2.0. The reproducibility supplement records the package versions, loaded **ggcircular** source and rendering session information needed to audit the computational environment.

Generative AI Statement

During the preparation of this article, the author used OpenAI Codex to assist with reviewing the accompanying code and checking the reproducibility of the computational workflow. Codex also assisted with typographic adjustments to figure labels during copyediting. OpenAI GPT-5.6 was used to assist with editorial revision of the manuscript, focusing on language, clarity, and readability. These tools were used in a supporting capacity. The author remains fully responsible for the manuscript and accompanying code, including the accuracy of the analyses and results, the interpretation of the findings, the originality of the work, and the proper attribution of all sources.

References

- Agostinelli, C. and Lund, U. (2025). **circular**: *Circular Statistics*, DOI: [10.32614/CRAN.package.circular](https://doi.org/10.32614/CRAN.package.circular). R package version 0.5-2.
- Brunson, J. C. (2020). **ggalluvial**: Layered grammar for alluvial plots. *Journal of Open Source Software*, 5(49):2017, DOI: [10.21105/joss.02017](https://doi.org/10.21105/joss.02017).
- Buja, A., Cook, D., Hofmann, H., Lawrence, M., Lee, E.-K., Swayne, D. F., and Wickham, H. (2009). Statistical inference for exploratory data analysis and model diagnostics. *Philosophical Transactions of the Royal Society A*, 367(1906):4361–4383, DOI: [10.1098/rsta.2009.0120](https://doi.org/10.1098/rsta.2009.0120).

- Di Marzio, M., Panzera, A., and Taylor, C. C. (2009). Local polynomial regression for circular predictors. *Statistics & Probability Letters*, 79(19):2066–2075, DOI: [10.1016/j.spl.2009.06.014](https://doi.org/10.1016/j.spl.2009.06.014).
- Di Marzio, M., Panzera, A., and Taylor, C. C. (2011). Kernel density estimation on the torus. *Journal of Statistical Planning and Inference*, 141(6):2156–2173, DOI: [10.1016/j.jspi.2011.01.002](https://doi.org/10.1016/j.jspi.2011.01.002).
- Di Marzio, M., Panzera, A., and Taylor, C. C. (2013). Non-parametric regression for circular responses. *Scandinavian Journal of Statistics*, 40(2):238–255, DOI: [10.1111/j.1467-9469.2012.00809.x](https://doi.org/10.1111/j.1467-9469.2012.00809.x).
- Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall/CRC, New York, DOI: [10.1201/9780429246593](https://doi.org/10.1201/9780429246593).
- García-Portugués, E., Crujeiras, R. M., and González-Manteiga, W. (2013). Kernel density estimation for directional-linear data. *Journal of Multivariate Analysis*, 121:152–175, DOI: [10.1016/j.jmva.2013.06.009](https://doi.org/10.1016/j.jmva.2013.06.009).
- García-Portugués, E. and Prieto-Tirado, A. (2023). Toroidal PCA via density ridges. *Statistics and Computing*, 33(5):107, DOI: [10.1007/s11222-023-10273-9](https://doi.org/10.1007/s11222-023-10273-9).
- Hong, S. and Jung, S. (2022). ClusTorus: An R package for prediction and clustering on the torus by conformal prediction. *The R Journal*, 14(2):186–207, DOI: [10.32614/RJ-2022-032](https://doi.org/10.32614/RJ-2022-032).
- Jammalamadaka, S. R. and SenGupta, A. (2001). *Topics in Circular Statistics*. World Scientific, Singapore, DOI: [10.1142/4031](https://doi.org/10.1142/4031).
- Mardia, K. V. and Jupp, P. E. (2000). *Directional Statistics*. John Wiley & Sons, Chichester, DOI: [10.1002/9780470316979](https://doi.org/10.1002/9780470316979).
- National Data Buoy Center (2024). Station 44013 historical standard meteorological data, 2024. <https://www.ndbc.noaa.gov/data/historical/stdmet/44013h2024.txt.gz>. Accessed 2026-06-18.
- National Data Buoy Center (2026a). Measurement descriptions and units. <https://www.ndbc.noaa.gov/faq/measdes.shtml>. Accessed 2026-06-18.
- National Data Buoy Center (2026b). Station 44013: Boston approach lighted buoy bf noaa 44013. https://www.ndbc.noaa.gov/station_page.php?station=44013. Accessed 2026-06-18.
- Nicosia, A. (2026). **ggcircular**: A **ggplot2** Extension for Circular and Directional Data, <https://github.com/AurelienNicosiaULaval/ggcircular>. R package development version 0.2.0 used for manuscript rendering.
- Pewsey, A. and García-Portugués, E. (2021). Recent advances in directional statistics. *TEST*, 30(1):1–58, DOI: [10.1007/s11749-021-00759-x](https://doi.org/10.1007/s11749-021-00759-x).

- Riehmman, P., Hanfler, M., and Froehlich, B. (2005). Interactive Sankey diagrams. In *IEEE Symposium on Information Visualization*, pages 233–240. DOI: [10.1109/INFVIS.2005.1532152](https://doi.org/10.1109/INFVIS.2005.1532152).
- Rosvall, M. and Bergstrom, C. T. (2010). Mapping change in large networks. *PLoS ONE*, 5(1):e8694, DOI: [10.1371/journal.pone.0008694](https://doi.org/10.1371/journal.pone.0008694).
- Wand, M. P. and Jones, M. C. (1995). *Kernel Smoothing*. Chapman and Hall, London, ISBN: [9780412552700](https://doi.org/10.1007/9780412552700).
- Wickham, H. (2010). A layered grammar of graphics. *Journal of Computational and Graphical Statistics*, 19(1):3–28, DOI: [10.1198/jcgs.2009.07098](https://doi.org/10.1198/jcgs.2009.07098).
- Wilkinson, L. (2005). *The Grammar of Graphics*. Springer, New York, 2nd edition, DOI: [10.1007/0-387-28695-0](https://doi.org/10.1007/0-387-28695-0).
- Xu, D. and Wang, Y. (2020). Area-proportional visualization for circular data. *Journal of Computational and Graphical Statistics*, 29(2):351–357, DOI: [10.1080/10618600.2019.1654881](https://doi.org/10.1080/10618600.2019.1654881).

A. Appendix: Supplementary Method and Displays

The companion displays complete the implemented grammar. The phase loom and toroidal flow represent conditional sector masses. The statistical orbit summarizes $\hat{\mu}_X(\theta)$ and local spread when a low-dimensional summary is scientifically appropriate. The flow displays are related to alluvial and Sankey visualizations ([Riehmman et al. 2005](#); [Rosvall and Bergstrom 2010](#); [Brunson 2020](#)), but their statistical target here is conditional circular sector mass.

For the statistical orbit, the conditional mean and local dispersion are

$$\hat{\mu}_X(\theta) = \frac{\sum_i K_\kappa(\theta - \theta_i) x_i}{\sum_i K_\kappa(\theta - \theta_i)}$$

and

$$\hat{\sigma}_X^2(\theta) = \frac{\sum_i K_\kappa(\theta - \theta_i) (x_i - \hat{\mu}_X(\theta))^2}{\sum_i K_\kappa(\theta - \theta_i)}.$$

The estimator $\hat{\mu}_X(\theta)$ is a circular-predictor analogue of a Nadaraya-Watson local constant smoother with circular kernel weights.

Experimental three-dimensional companions are implemented in the software for geometric inspection of cylinder and torus surfaces. They are optional aids and are not used as evidence in the article.




Figure A1: Phase Loom on a smooth cylindrical scenario. Flows represent conditional masses $P(X \in B_j \mid \Theta \in A_k)$.




Figure A2: Statistical Orbit for a heteroscedastic cylindrical scenario. The central curve estimates $\hat{\mu}_X(\theta)$ and the ribbon represents local spread.













Figure A3: Toroidal Flow for a diagonal scenario. Conditional masses move toward sectors of Φ aligned with Θ .

Table A1: Detailed visual positioning of the implemented graphical grammar relative to existing displays and estimands.

View	Existing display	Proposed display	Main limitation	Proposed advantage
Joint KDE, $\hat{f}(\theta, x)$			Joint relief mixes conditional shape with the marginal density of Θ .	Row normalization targets $\hat{f}(x \theta)$ and makes angular slices comparable.
Circular-predictor regression			A single mean trend can hide multimodality and changing spread.	Topography and orbit separate full shape from low-dimensional summaries.
Toroidal density, $\hat{f}(\theta, \phi)$			Joint density does not organize reading by conditioning angle θ .	Toroidal topography targets $\hat{f}(\phi \theta)$ directly.
Toroidal ridge			A ridge-only summary can detach the curve from surrounding mass.	The ridge layer ties $\hat{r}(\theta)$ to density and concentration.
Alluvial flow			Generic flows are not periodic and not conditional by default.	Phase loom and toroidal flow attach masses to circular conditioning.
Marginal circular graphics			Marginal summaries can miss dependence with weak marginals.	Conditional displays target $S^1 \times \mathbb{R}$ and $S^1 \times S^1$ dependence.

The simulation protocol uses the following scenarios.

Table A2: Simulation scenarios used to evaluate the graphical diagnostics. Here $\text{VM}(\mu, \kappa)$ denotes a von Mises distribution with mean direction μ and concentration κ .

Space	Scenario	Data-generating mechanism
Cylindrical	independent	$\Theta \sim U(0, 2\pi)$ and $X \sim N(3, 0.75^2)$
Cylindrical	smooth	$X = 3 + 0.95 \cos(\Theta - 0.35) + 0.35 \sin(2\Theta + 0.5) + \epsilon$, $\epsilon \sim N(0, 0.28^2)$
Cylindrical	heteroscedastic	$X = 3 + \epsilon(\Theta)$ with $\text{sd}\{\epsilon(\Theta)\} = 0.20 + 0.95\{1 + \sin(\Theta - 0.8)\}/2$
Cylindrical	multimodal	Equal mixture of two conditional Gaussian branches with means $2.55 + 0.55 \cos(\Theta - 0.3)$ and $3.75 - 0.60 \sin(\Theta + 0.6)$, each with $\epsilon \sim N(0, 0.18^2)$
Cylindrical	seam	$X = 3 + 1.05 \cos(\Theta) + \epsilon$, $\epsilon \sim N(0, 0.25^2)$
Toroidal	independent concentrated	$\Phi \sim \text{VM}(1.2, 7)$ independently of Θ
Toroidal	diagonal	$\Phi \sim \text{VM}(\Theta, 12)$
Toroidal	doubling	$\Phi \sim \text{VM}(2\Theta, 12)$
Toroidal	bifurcation	Equal mixture of $\text{VM}\{\Theta + 0.65 \sin(\Theta), 14\}$ and $\text{VM}\{\Theta + \pi - 0.65 \sin(\Theta), 14\}$
Toroidal	uniform-marginal dependence	Equal antipodal mixture of $\text{VM}(\Theta, 18)$ and $\text{VM}(\Theta + \pi, 18)$

The diagnostic summaries are:

Table A3: Scalar diagnostics used in the simulation study. They are descriptive summaries of the displays, not replacements for visual inspection.

Display	Diagnostic statistic	Interpretation
Phase Loom	Cramér's V after binning Θ into 24 sectors and X into 16 quantile bins	Conditional mass variation across angular sectors
Circular Topography	Mean squared deviation of row-normalized density rows from their angular average	Variation in $\hat{f}(x \theta)$
Statistical Orbit	Scaled range of $\hat{\mu}_X(\theta)$ and $\hat{\sigma}_X(\theta)$	Variation in local mean or spread
Toroidal Flow	Cramér's V after binning both angles into 24 sectors	Conditional angular mass variation
Toroidal Topography	Mean squared deviation of row-normalized density rows from their angular average	Variation in $\hat{f}(\phi \theta)$
Conditional Ridge	Resultant-weighted circular variation of $\hat{r}(\theta)$	Nonconstant modal relation on the torus

Figure A4 compares the proposed conditional topography with the joint directional-linear KDE that underlies classical density displays.




Figure A4: Comparison with directional-linear KDE. Left: joint KDE on $S^1 \times \mathbb{R}$. Right: conditional topography using the same kernel estimates after row normalization.

For circular-linear data, the smoothing selection score is

$$CV(\kappa, h) = \frac{1}{n} \sum_{i=1}^n \log \hat{f}_{-v(i), \kappa, h}(x_i | \theta_i),$$

Table A4: NDBC variables used in the empirical illustration.

Variable	Meaning	Unit	Role
WDIR	wind direction	degrees true	conditioning angle
WSPD	wind speed	m/s	linear response in the cylindrical panel
MWD	mean wave direction	degrees true	circular response in the toroidal panel
WVHT	wave height	m	preprocessing context only

Table A5: Preprocessing counts for the NDBC empirical illustration.

Step	Count	Rule
Raw records	52,665	records read from the NDBC file
Valid wind direction	52,640	$0 \leq \text{degrees} < 360$; code 999 removed
Valid wind speed	52,640	m/s ≥ 0 ; code 99 removed
Valid wave direction	17,108	$0 \leq \text{degrees} < 360$; code 999 removed
Valid wave height	17,424	m ≥ 0 ; code 99 removed
Final wind-speed pairs	52,640	valid WDIR and WSPD
Final wind-wave pairs	17,106	valid WDIR, WSPD and MWD

where $v(i)$ denotes the validation fold and $\hat{f}_{-v(i)}$ is fitted without the observations in that fold. The five folds are constructed by a balanced repetition of the labels 1 to 5 followed by a seeded random permutation, so fold sizes differ by at most one. For toroidal data, the same criterion is applied to $\hat{f}(\phi_i \mid \theta_i)$ over candidate pairs $(\kappa_\theta, \kappa_\phi)$. The response-angle kernel includes the von Mises normalizing constant so that log scores are comparable across κ_ϕ values. The generated reproduction tables record the tuning-row construction, fold seeds, candidate grids and selected values.

Table A6: Generated smoothing-selection protocol for the NDBC displays. Tuning rows are complete records selected at equally spaced positions in the preprocessed data.

Space	Tuning rows	Folds	Fold seed	κ candidates	h/h_0 candidates
Cylindrical	5000	5	20260615	8;14;24;36	0.6;0.8;1;1.3;1.7
Toroidal	5000	5	20260616	8;14;24;36	—

B. Appendix: Diagnostic and Supplementary Figures

Figure A5 summarizes the two sample spaces used in the article. The cylindrical case combines one periodic coordinate and one real coordinate. The toroidal case identifies the edges of the unfolded rectangle in both angular directions.


Figure A5: Sample spaces $S^1 \times \mathbb{R}$ and $S^1 \times S^1$.

Figure A6 illustrates descriptive bootstrap intervals for a smoothed conditional mean and for local concentration along a toroidal ridge (Efron and Tibshirani 1993). The bootstrap is used here as a visualization of empirical variability, not as a formal test.




Figure A6: Bootstrap uncertainty for two conditional summaries. Left: smoothed conditional mean. Right: local concentration along a toroidal ridge.

Figure A7 reports pointwise, not simultaneous, simulation coverage diagnostics corresponding to the numerical summaries in the main text. Each cell uses 50 independently generated datasets, 200 bootstrap resamples per dataset and 45 unique conditioning directions. The cylindrical plotting function appends 2π only to close its displayed curve, and this graphical copy of direction 0 is removed before coverage is evaluated. The toroidal grid instead consists directly of 45 distinct cell centres and has no appended endpoint. Thus each cell has 2250 evaluations when all quantities are finite. The output files retain covered and evaluable counts, coverage and a Monte Carlo standard error across the 50 dataset-level proportions. Toroidal 95% intervals use the 0.95 quantile of circular bootstrap distances from their circular centre and target the nearest generating modal branch; cylindrical percentile intervals target the generating conditional

mean. The minimum coverage is 0.506 for the cylindrical smooth scenario at $n = 2000$ and $\kappa = 14$, consistent with smoothing bias because the intervals target a smoothed curve without bias correction for the unsmoothed generating truth. For the bifurcation scenario, coverages from 0.972 to 0.993 accompany mean circular interval widths from 1.71 to 3.42 radians. These wide intervals can cover either nearby generating branch and therefore do not indicate good calibration despite their high coverage.




Figure A7: Pointwise bootstrap coverage across the simulation protocol. Rows separate $\kappa = 14$ and $\kappa = 24$, and all 30 values are printed in their tiles. Each cell pools 50 datasets over 45 unique conditioning directions; the intervals are not simultaneous bands.

Figures A8 and A9 give the full display-by-scenario detection grid. Displays are arranged in rows and the four smoothing values are separated into columns. Each map contains 240 distinct cells, so the 480 protocol cells appear exactly once. The main text uses the more compact detection summary because these full grids are audit-oriented.




Figure A8: Detailed cylindrical simulation detection rates. Displays are shown in rows and $\kappa \in \{8, 14, 24, 36\}$ in columns. Exact rates are provided in `output/tables/simulation_detection_rates.csv`.




Figure A9: Detailed toroidal simulation detection rates. Displays are shown in rows and $\kappa \in \{8, 14, 24, 36\}$ in columns. Exact rates are provided in `output/tables/simulation_detection_rates.csv`.

Figure A10 documents two methodological diagnostics. The left panel shows compact sensitivity summaries as κ changes. The right panel shows an example of low marginal support for the conditioning angle.




Figure A10: Smoothing sensitivity and marginal support diagnostics.

Figure A11 reports timing summaries for the main data preparation functions on moderate fixed grids.




Figure A11: Computation times for the main graph-preparation functions.

C. Appendix: Empirical Sensitivity

The reproducibility supplement contains the fixed-seed scripts and generated outputs needed to audit the numerical results. Detailed reproduction commands are documented in the accompanying README rather than repeated in the manuscript. The table below reports the sensitivity of the empirical permutation diagnostics to angular discretization.

Table A7: Permutation sensitivity to angular discretization. Each response is conditioned on wind direction. With 999 permutations, 0.001 is the smallest attainable p-value under the add-one convention.

Response	n	Theta bins	Response bins	V	p	Null mean	Null 95%
Wind speed	52,640	12	10	0.134	0.001	0.014	0.016
Wave direction	17,106	12	12	0.275	0.001	0.025	0.028
Wind speed	52,640	24	20	0.102	0.001	0.021	0.022
Wave direction	17,106	24	24	0.215	0.001	0.037	0.038
Wind speed	52,640	36	20	0.106	0.001	0.026	0.027
Wave direction	17,106	36	36	0.183	0.001	0.045	0.047

All rows in this sensitivity table use exchangeable, 999 permutations.

Affiliation:

Aurélien Nicosia
Department of Mathematics and Statistics
Université Laval
Québec City, Québec, Canada
E-mail: aurelien.nicosia@mat.ulaval.ca