Localized Cluster Enhancement: TFCE revisited with valid error control

Samuel Davenport, Wouter Weeda, Thomas E. Nichols, Jelle Goeman

University of California, San Diego

March 3, 2025

Notation

- Let $\mathcal{B}$ denote the set of voxels/vertices making up the brain.
- Suppose that, for each $v \in \mathcal{B}$, X_v are the underlying data corresponding to voxel v. (Containing data from multiple subjects.)
- Suppose that at each v we have a test statistic $T_v = T(X_v)$.

TFCE recap

- TFCE is a widely used method for identifying activation with over 5000 citations.
- At each voxel v, TFCE transforms T to

$$S_v(T) = \int_{h_0}^{T_v} h^H e_v(h, T)^E \mathrm{d}h.$$

We can also write $S_v(T) = S_v(\mathbf{X})$
- Here h is the height and $e_v(h, T)$ is the extent of the test statistic at the level h for the voxel v.

Understanding the TFCE integral

$$S_v(T) = \int_{h_0}^{T_v} h^H e_v(h,T)^E \, dh$$

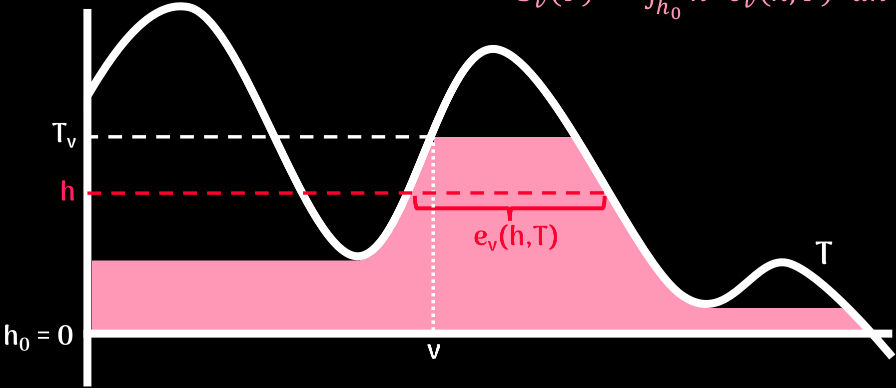

TFCE recap

$$S_v(T) = \int_{h_0}^{T_v} h^H e_v(h, T)^E \mathrm{d}h.$$

- In practice - $H = 2$, $E = 0.5$ and $h_0 = 0$ are the default parameters typically chosen.
- One of the main reasons for introducing
- Permuted TFCE test-statistics: $S^*_{v,1}, \dots, S^*_{v,P}$ are calculated.
- Given $\alpha \in (0, 1)$, a TFCE threshold t^* is chosen based on the $(1 - \alpha)\%$ quantile of the permutation distribution of $\max_{1 \leq p \leq P} S^*_{v,p}$. Reject v such that $S_v(T) > t^*$.

Approximating TFCE

In practice the integral is computed numerically. As such we in fact approximate

$$S_v(T) = \int_{h_0}^{T_v} h^H e_v(h, T)^E \mathrm{d}h.$$

by

$$S_v(T) = \delta \sum_{j=0}^{\lfloor (T_v - h_0)/\delta \rfloor} (h_0 + j\delta)^H e_v(h_0 + j\delta)^E. \quad (1)$$

Note that this introduces a 4th arbitrary δ parameter into the mix, one which can also affect the outcome. FSL uses $\delta = 0.1$ but this is arbitrary.

TFCE claims

Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference

Stephen M. Smith[a,*] and Thomas E. Nichols[b,a,c]

Many image enhancement and thresholding techniques make use of spatial neighbourhood information to boost belief in extended areas of signal. The most common such approach in neuroimaging is cluster-based thresholding, which is often more sensitive than voxel-wise thresholding. However, a limitation is the need to define the initial cluster-forming threshold. This threshold is arbitrary, and yet its exact choice can have a large impact on the results, particularly at the lower (e.g., t, z < 4) cluster-forming thresholds frequently used. Furthermore, the amount of spatial pre-smoothing is also arbitrary (given that the expected signal extent is very rarely known in advance of the analysis). In the light of such problems, we propose a new method which attempts to keep the sensitivity benefits of cluster-based thresholding (and indeed the general concept of "clusters" of signal), while avoiding (or at least minimising) these problems. The method takes a raw statistic image and

holding. For inference, the TFCE image can easily be turned into voxel-wise *p*-values (either uncorrected, or corrected for multiple comparisons across space) via permutation testing.

Unfortunately TFCE does not address these problems and in fact greatly exacerbates them.

Thresholded TFCE

We apply TFCE to 20 subjects from the HCP to the primary social contrast.

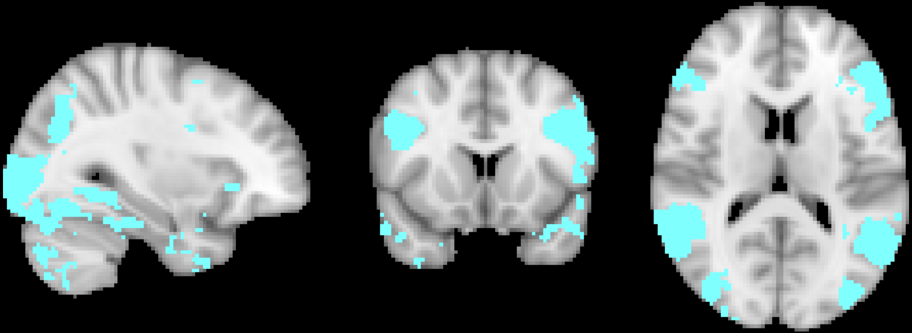

However TFCE borrows information from across its support.

TFCE support

We apply TFCE to 20 subjects from the HCP to the primary social contrast.

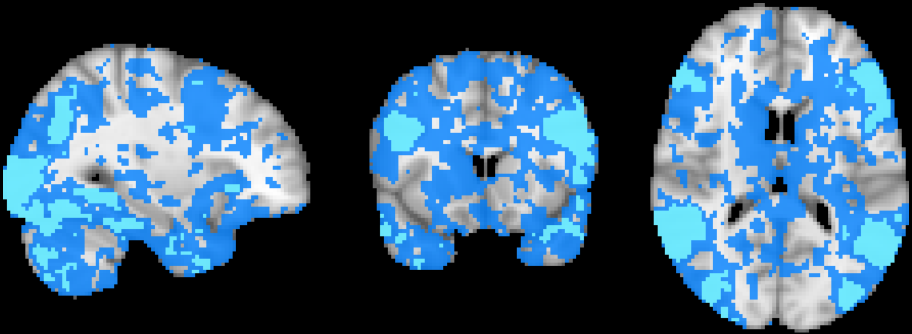

TFCE thus can be used to make a global statement but stuggles to localize activation.

Simulation example (based on $n = 50$)

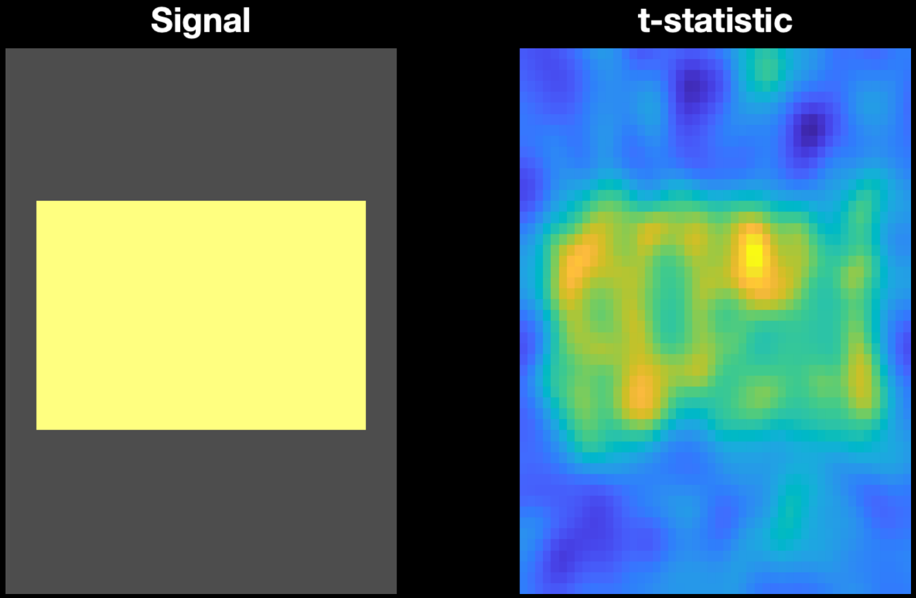

Simulation example (based on $n = 50$)

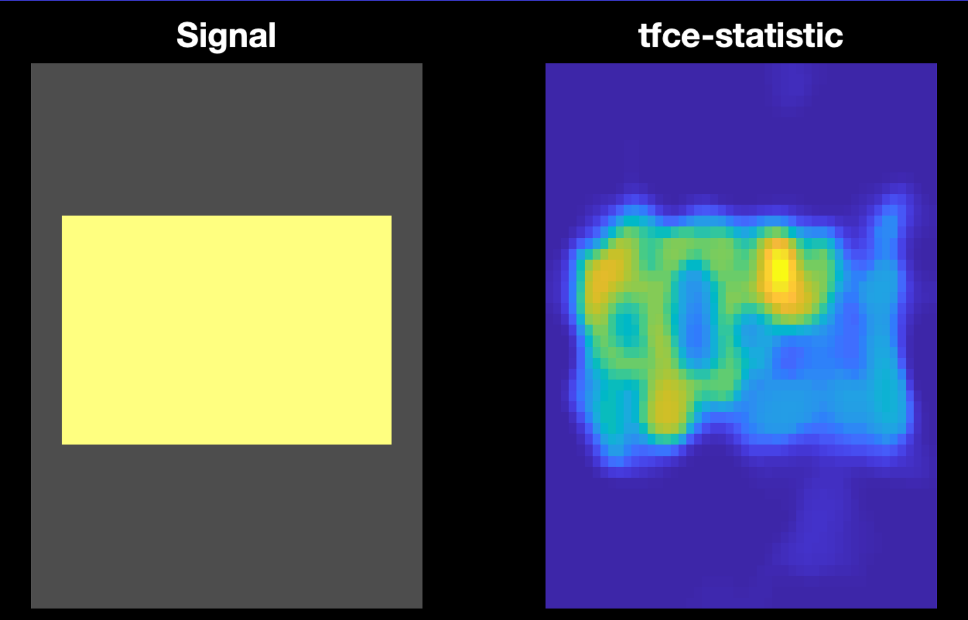

Simulation example (based on $n = 50$)

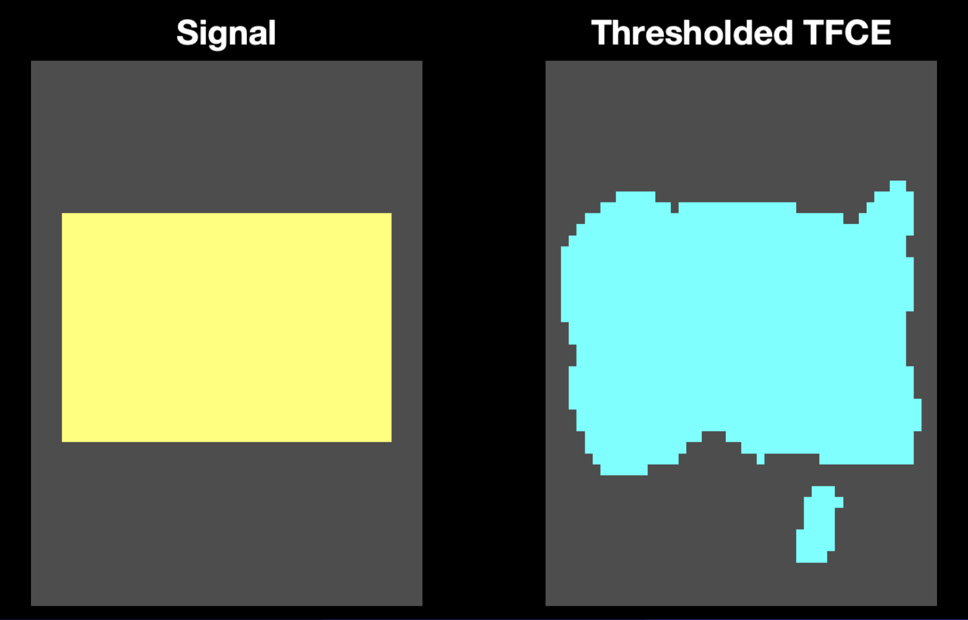

Simulation example (based on $n = 50$)

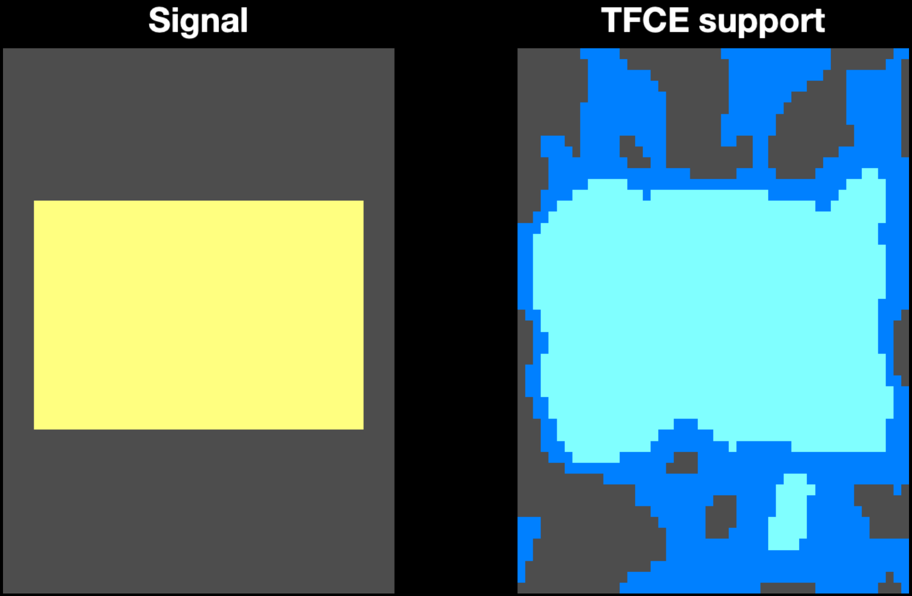

Simulation example (based on $n = 50$)

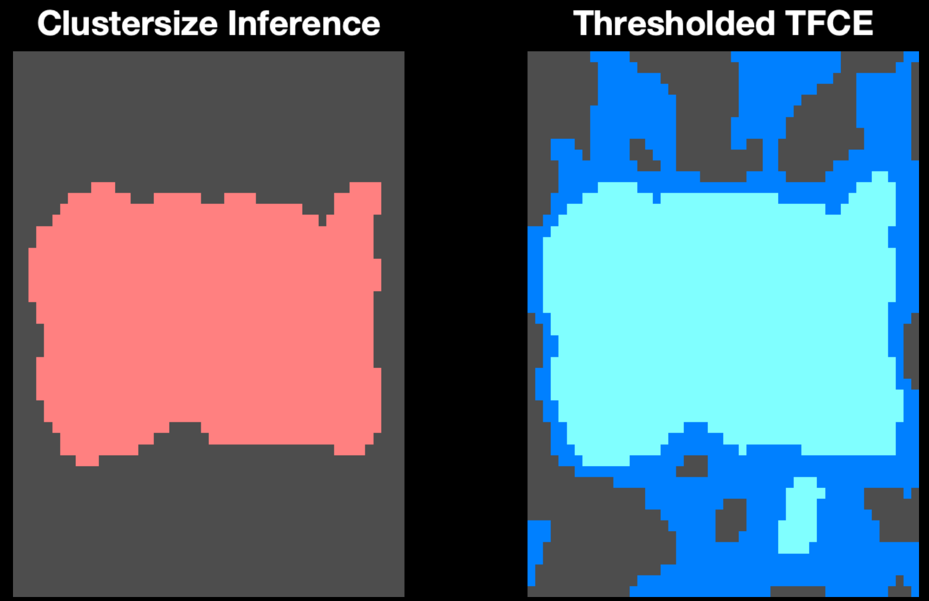

TFCE can have inflated error rates

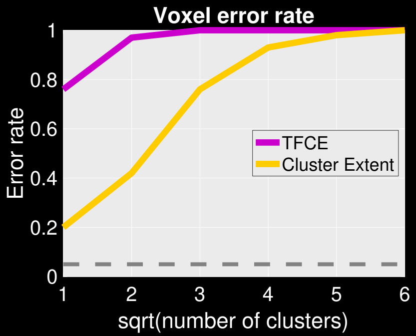

TFCE can have inflated error rates

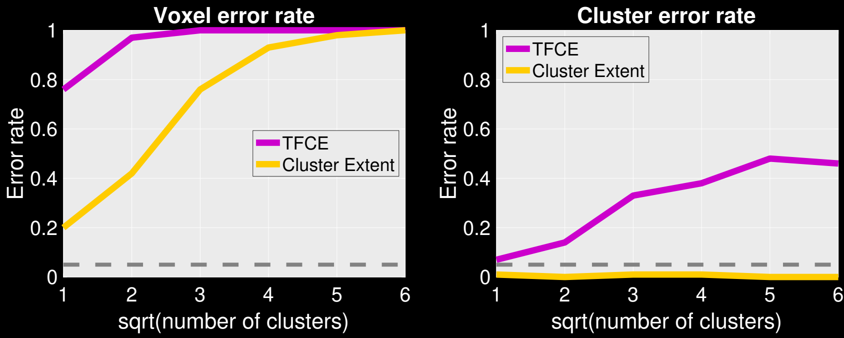

Global vs voxel vs cluster level inference

There are 3 types of inference statements typically used.

1. Voxel: Every highlighted voxel is active
2. Cluster: Every cluster contains at least one active voxel
3. Global: There is some voxel active somewhere in the brain.

Classifying fMRI inference methods

There are 3 types of inference statements typically used.

1. Voxel: Every highlighted voxel is active
Voxelwise inference

2. Cluster: Every highlighted cluster contains at least one active voxel
Clustersize inference

3. Global: There is some voxel active somewhere in the brain.
TFCE

TFCE issues

- TFCE was designed to improve on cluster extent inference however it doesn't control cluster error rates
- One of the main motatavations for TFCE was to reduce researcher degrees of freedom, to make it "threshold free". However the TFCE transformation is defined as:

$$S_v = \int_{h_0}^{T_v} h^H e_v(h)^E \mathrm{d}h. \tag{2}$$

- In particular h_0 acts as a threshold so TFCE is not threshold free.

Localized Cluster Enhancement

To localize TFCE for a region $R \subset \mathcal{V}$ let

$$S_{v,R}(T) = S_v(T \times 1[R]) = \int_{h_0}^{T_v} h^H e_v(h, T \times 1[R])^E \mathrm{d}h.$$

Then we can say R contains at least one active voxel if $\max_{v \in R} S_{v,R}(T) > t^*$ where t^* is the original TFCE cutoff.

Generalized TFCE statistic

$$S_v = \int_{h_0}^{\infty} f(h) g\big(e_v(h)\big) \mathrm{d}h. \tag{3}$$

- This reduces to the TFCE-statistic if we take $f(x) = x^H$ and $g(x) = x^E$, and note that $e_v(h) = 0$ if $h > T_v$.
-We recover **cluster-size inference** by taking f to be a Dirac δ-function shifted to h_0, and g the identity.
- **Cluster-mass inference** can be obained by taking $f(x) = 1$ and g the identity
- and **peak height inference** with f the identity and $g(x)$ the indicator of $x > 0$.

Exchangeability Assumption

Let $\mathcal{I} \subseteq \mathcal{B}$ be the collection of all inactive voxels, and Π be the group of permutations considered.
We shall require the following **exchangeability assumption**.

Assumption

For every $\pi \in \Pi$, we have $(\mathbf{X}_v)_{v\in\mathcal{I}} \overset{d}{=} (\mathbf{X}_v \circ \pi)_{v\in\mathcal{I}}$.

Moreover let $H_{\mathcal{B}}$ be the null hypothesis that all voxels in $\mathcal{B}$ are inactive.

TFCE provides weak FWER control

The TFCE threshold t^* is the kth order statistic of

$$\max_{v \in \mathcal{B}} S_v(\mathbf{X} \circ \pi_1), \dots, \max_{v \in \mathcal{B}} S_v(\mathbf{X} \circ \pi_P),$$

where $k = \lceil (1-\alpha)P \rceil$.

Theorem

Assume that Exchangability holds. Suppose that π_1 is the identity permutation and $\pi_2, \dots, \pi_P$ are drawn independent and uniformly, with replacement, from Π. Then, if $H_{\mathcal{B}}$ is true,

$$\mathrm{P}\Big(\max_{v \in \mathcal{B}} S_v(\mathbf{X}) > t^*\Big) \leq \alpha.$$

Constructing a local teset

Given $\mathcal{R} \subseteq \mathcal{B}$, consider the masked data $\mathbf{X}_\mathcal{R} = (X_v 1[v \in \mathcal{R}])_{v\in\mathcal{B}}$. Let $H_\mathcal{R}$ be the null hypothesis that all voxels in $\mathcal{R}$ are inactive. Let $t^*_\mathcal{R}$ be the kth order statistic of

$$\max_{v\in\mathcal{R}} S_v(\mathbf{X}_\mathcal{R} \circ \pi_1), \dots, \max_{v\in\mathcal{B}} S_v(\mathbf{X}_\mathcal{R} \circ \pi_P),$$

where $\mathbf{X}_\mathcal{R} \circ \pi = (X_v \circ \pi)_{v\in\mathcal{R}}$.

Theorem

Assume that Exchangability holds. Suppose that π_1 is the identity permutation and $\pi_2, \dots, \pi_P$ are drawn independent and uniformly, with replacement, from Π. Then, if $H_\mathcal{R}$ is true,

$$\mathrm{P}\Big(\max_{v\in\mathcal{R}} S_v(\mathbf{X}_\mathcal{R}) > t^*_\mathcal{R}\Big) \leq \alpha.$$

Localized Cluster Enhancement

In order to provide simultaneously valid inference over all regions $\mathcal{R} \subset \mathcal{B}$ we introduce Localized Cluster Enhancement.

$$\mathrm{LCE}(\mathcal{R}) = \frac{1}{P}\sum_{p=1}^{P} 1\left[\max_{v\in\mathcal{R}} S_v(\mathbf{X}_{\mathcal{R}}) \leq \max_{v\in\mathcal{B}} S_v(\mathbf{X}\circ\pi_p)\right]. \quad (4)$$

We then reject $H_{\mathcal{R}}$ for a region $\mathcal{R} \subseteq \mathcal{B}$ if $\mathrm{LCE}(\mathcal{R}) \leq \alpha$.

Simultaneous validity of LCE

Recall
$\text{LCE}(\mathcal{R}) = \frac{1}{P}\sum_{p=1}^{P} 1\left[\max_{v\in\mathcal{R}} S_v(\mathbf{X}_{\mathcal{R}}) \leq \max_{v\in\mathcal{B}} S_v(\mathbf{X}\circ\pi_p)\right].$

Theorem

(Simultaneous validity of LCE) Assume that Exchangability holds. Suppose that π_1 *is the identity permutation and* $\pi_2, \dots, \pi_P$ *are drawn independent and uniformly, with replacement, from* Π*. Given* $\alpha \in (0,1)$*, we have*

$$\mathbb{P}\Big(\textit{there exists } \mathcal{R} \subseteq \mathcal{B} \textit{ such that } H_{\mathcal{R}} \textit{ is true and } LCE(\mathcal{R}) \leq \alpha\Big) \leq \alpha.$$

Equivalently $\mathbb{P}\left(LCE(\mathcal{R}) \geq \alpha \;\; \textit{for all } \mathcal{R} \subseteq \mathcal{I}\right) \geq 1-\alpha.$

Characterizing LCE

We can characterize LCE rejected regions using the original TFCE threshold.

Corollary

Let t^* *be the TFCE threshold. Then LCE rejects a region* $\mathcal{R} \subseteq \mathcal{B}$ *if and only if* $\max_{v \in \mathcal{R}} S_v(\mathbf{X}_\mathcal{R}) > t^*$. *As such, under the same assumptions as before*

$$\mathbb{P}\Big(\textit{there exists } \mathcal{R} \subseteq \mathcal{B} \textit{ such that } H_\mathcal{R} \textit{ is true and } \max_{v \in \mathcal{R}} S_v(\mathbf{X}_\mathcal{R}) > t^*\Big)$$

Localized Cluster Enhancement illustration

The t-statistic based on 20 subjects (for the HCP Gambling task)

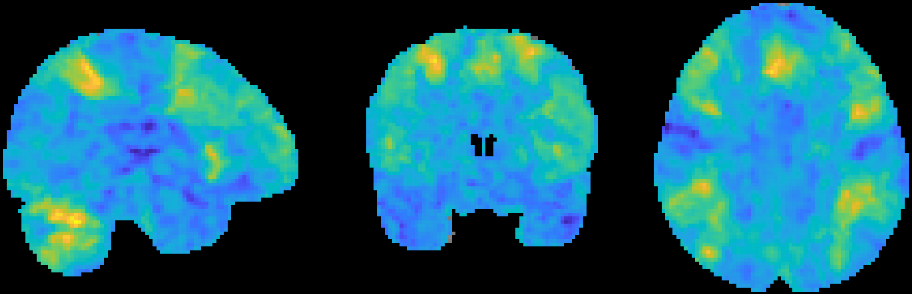

Localized Cluster Enhancement illustration

Apply a mask of the Middle Frontal Gyrus

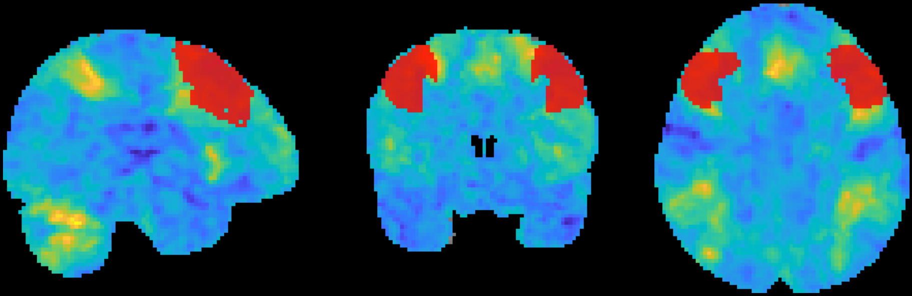

Localized Cluster Enhancement illustration

Apply a mask of the Middle Frontal Gyrus

Localized Cluster Enhancement illustration

Apply a mask of the Middle Frontal Gyrus

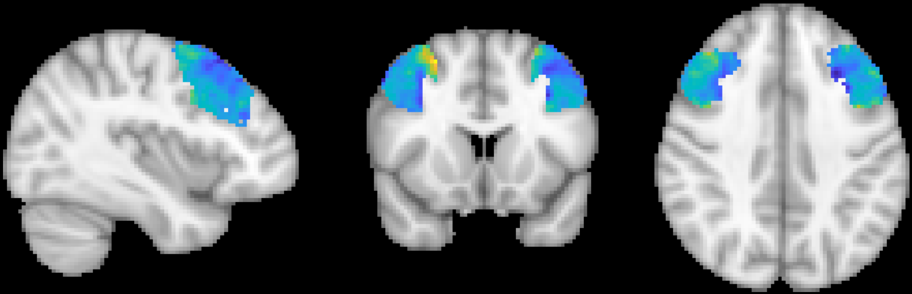

Localized Cluster Enhancement illustration

Apply TFCE transformation to obtain $\max_{v \in R} S_v(\mathbf{X}_{\mathcal{R}})$

Reject if $\max_{v \in R} S_v(\mathbf{X}_{\mathcal{R}}) > t^*$.

LCE as an Algorithm

Algorithm 1 Localized Cluster Enhancement

Require: Data $\mathbf{X} = (X_v)_{v\in\mathcal{B}}$, real-valued functions f, g, a threshold $h_0 \in \mathbb{R}$, a set of n regions of interest $\{\mathcal{R}_i\}_{i=1}^n \subseteq \mathcal{B}$, a number of permutations P, and a collection of permutations Π.

Compute the original test-statistics $\{T_v\}_{v\in\mathcal{B}}$.

Compute the transformed statistics $S_v(\mathbf{X}) = \int_{h_0}^{\infty} f(h)g\big(e_v(h, \mathbf{X})\big)\mathrm{d}h$, for each $v \in \mathcal{B}$.

Draw permutations $\pi_1, \dots, \pi_P$ independently from Π.

for $p = 1$ to P **do**

  Compute the permuted transformed map: $S_v(\mathbf{X}\circ\pi_p) = \int_{h_0}^{\infty} f(h)g\big(e_v(h, \mathbf{X}\circ\pi_p)\big)\mathrm{d}h$, for $v \in \mathcal{B}$.

end for

for $i = 1$ to n **do**

  Mask the data to the region $\mathcal{R}_i$, to compute $\mathbf{X}_{\mathcal{R}_i} = (X_v 1[v \in \mathcal{R}_i])_{v\in\mathcal{B}}$.

  Compute the masked transformed map $S_v(\mathbf{X}_{\mathcal{R}_i}) = \int_{h_0}^{\infty} f(h)g\big(e_v(h, \mathbf{X}_{\mathcal{R}i})\big)\mathrm{d}h$.

  Calculate $\mathrm{LCE}(\mathcal{R}_i) = \frac{1}{P}\sum_{p=1}^{P} 1\left[\max_{v\in\mathcal{R}_i} S_v(\mathbf{X}_{\mathcal{R}_i}) \leq \max_{v\in\mathcal{B}} S_v(\mathbf{X}\circ\pi_p)\right]$.

end for

return $\{\mathrm{LCE}(\mathcal{R}_i)\}_{i=1}^n$: the set of LCE-adjusted p-values.

given $\alpha \in (0,1)$, reject $H_{\mathcal{R}_i}$ for each $1 \leq i \leq n$ such that $\mathrm{LCE}(\mathcal{R}_i) \leq \alpha$.

Advantages of Localized Cluster Enhancement

- If $\max_{v \in R} S_{v,R}(T) > t^*$, then LCE claims that there is at least one active voxel within the region R.
- This is provably valid (with an error rate of 5%) simultaneously over all regions R so LCE can make local claims unlike TFCE.
- LCE can be applied to pre-defined regions based on an atlas or to data-driven regions such as clusters.
- Allows the user to explore the data and find the regions R of interest.

Clusterwise rejections

LCE can be applied to the clusters obtained from TFCE to give rigorous guarantees.

Proposition

(Establishing the significance of clusters identified by TFCE) *Under the above assumptions, let* $\mathcal{C}_1, \dots, \mathcal{C}_m \subseteq \mathcal{B}$ *denote the clusters identified to be TFCE significant. Then*

$$\mathbb{P}\left(LCE(\mathcal{C}_i) \geq \alpha \text{ for all } 1 \leq i \leq n \text{ such that } \mathcal{C}_i \subseteq \mathcal{I}\right) \geq 1 - \alpha.$$

Characterizing the error control of TFCE

Corollary

Under above assumptions, let $\mathcal{C}_1, \ldots, \mathcal{C}_m \subseteq \mathcal{B}$ denote the clusters identified to be "TFCE significant". Then

$$\mathbb{P}\left(supp_{h_0}(\mathcal{C}_i) \subseteq \mathcal{I} \text{ for some } 1 \leq i \leq m\right) \leq \alpha.$$

This result really shows that h_0 effectively acts as a threshold - in the same way as the cluster forming threshold does for clustersize inference.

TFCE support

Going back to this image.

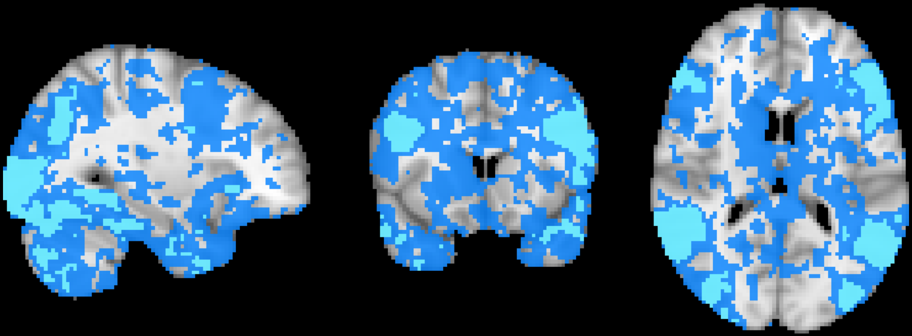

We have now formally shown that TFCE as classically used cannot localize activation.

Voxelwise statements using LCE

We can in fact make certain voxelwise statements using LCE. Note that given $v \in \mathcal{B}$, taking $\mathcal{R} = \{v\}$, the restricted TFCE statistic is

$$S_v(\mathbf{X}_{\{v\}}) = \int_{h_0}^{T_v} f(h)g(1)\, dh, \tag{5}$$

since $e_v(h) = 1[h_0 \leq h \leq T_v]$.
In the default TFCE setting, $f(h) = h^H$ and $g(1) = 1$, we have

$$S_v(\mathbf{X}_{\{v\}}) = \int_{h_0}^{T_v} h^H\, dh = \frac{1}{H+1}(T_v^{H+1} - h_0^{H+1}),$$

As such $T_v \geq (t^*(H+1) + h_0^{H+1})^{\frac{1}{H+1}}$ implies v is voxelwise significant with strong FWER conrol.

Voxelwise statements using LCE - cont

Recall: we have

$$S_v(\mathbf{X}_{\{v\}}) = \int_{h_0}^{T_v} h^H \, dh = \frac{1}{H+1}(T_v^{H+1} - h_0^{H+1}),$$

So when $H = 2$ and $h_0 = 0$, we have $\frac{H+1}{(}T_v^{H+1} - h_0^{H+1}) = \frac{T_v^3}{3}$. As such we may call all voxels such that $T_v > (3t^*)^{1/3}$ voxelwise significant and provide strong control of the FWER.
In practice this can be very high so is mainly of theoretical interest for $h_0 = 0$. However it may be more useful at higher levels of h_0.

Regional activation (HCP - Social)

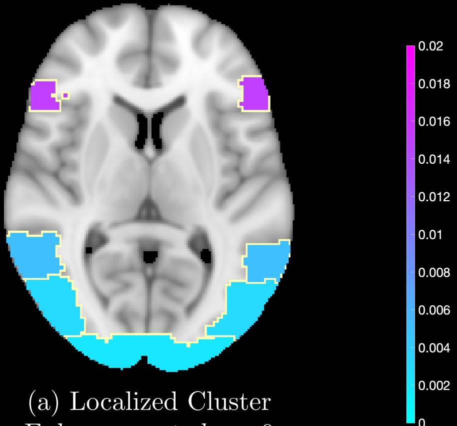

(a) Localized Cluster Enhancement, $h_0 = 0$

Regional activation (HCP - Social)

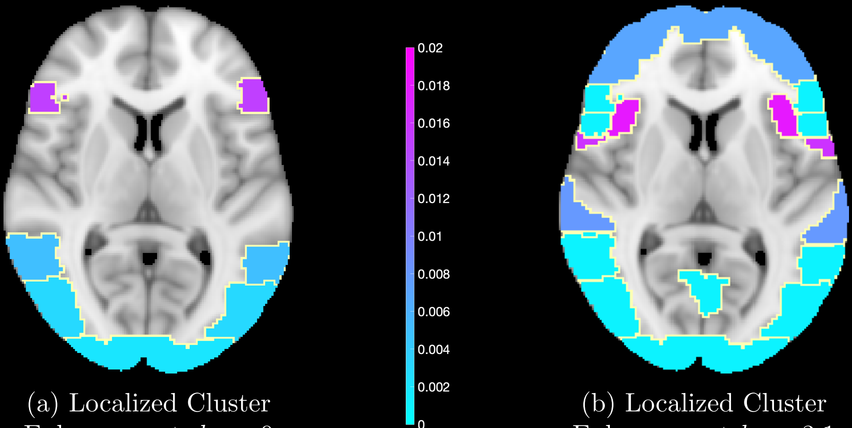

(a) Localized Cluster Enhancement, $h_0 = 0$

(b) Localized Cluster Enhancement $h_0 = 3.1$

Power on 20 subjects from the HCP

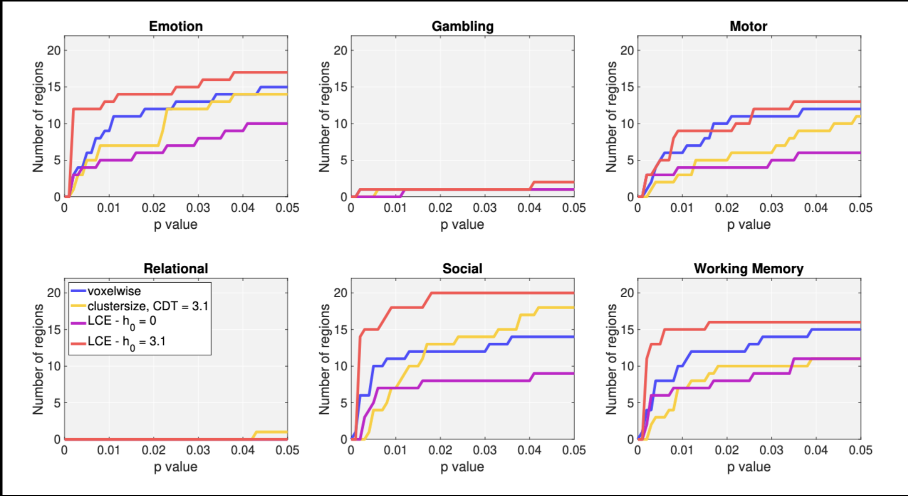

Power on 40 subjects from the HCP

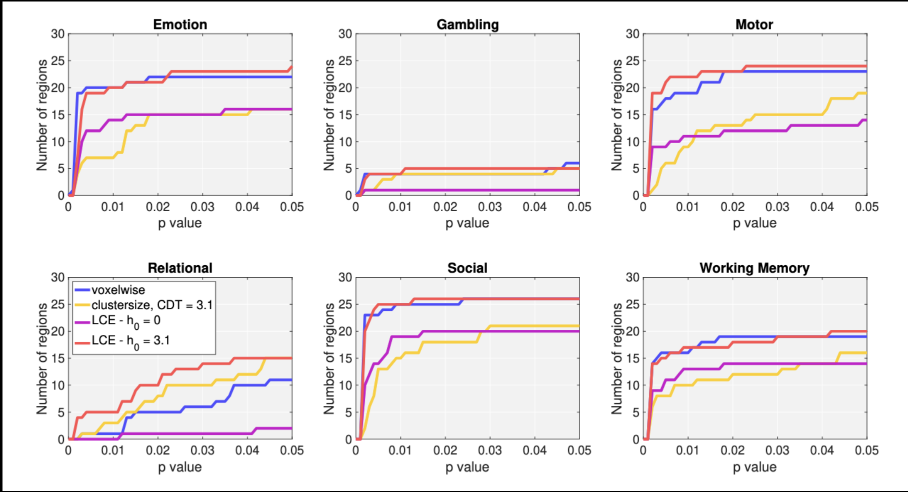

Classifying fMRI inference methods

There are 3 types of inference statements typically used.

1. Voxel: Every highlighted voxel is active
Voxelwise inference

2. Cluster: Every highlighted cluster contains at least one active voxel
LCE, Clustersize inference

3. Global: There is some voxel active somewhere in the brain.
TFCE

Conclusions

- TFCE as classically used can have inflated voxel and clusterwise error rates.

- **Localized Cluster Enhancement** provably controls clusterwise and regional error rates and allows for increases in power and localization.

- TFCE is not threshold free as it (strongly) depends on a threshold h_0. The default choice of $h_0 = 0$ means that TFCE typically can only make the weak global statement in practice.

- For localized cluster enchancement we recommend a threshold of $h_0 = 3.1$ in line with the default for clustersize inference.

Thanks

- Slides for this talk are available on my website:
`sjdavenport.github.io/talks`

- Code to implement LCE and a tutorial on TFCE are available in the StatBrainz MATLAB package available at:
`sjdavenport.github.io/software`

- Further details available at my poster: 1871