Multiple testing for modern data: structure, curation, and replicability

Eugene Katsevich

Department of Statistics
Stanford University

February 5, 2019

A modern data set
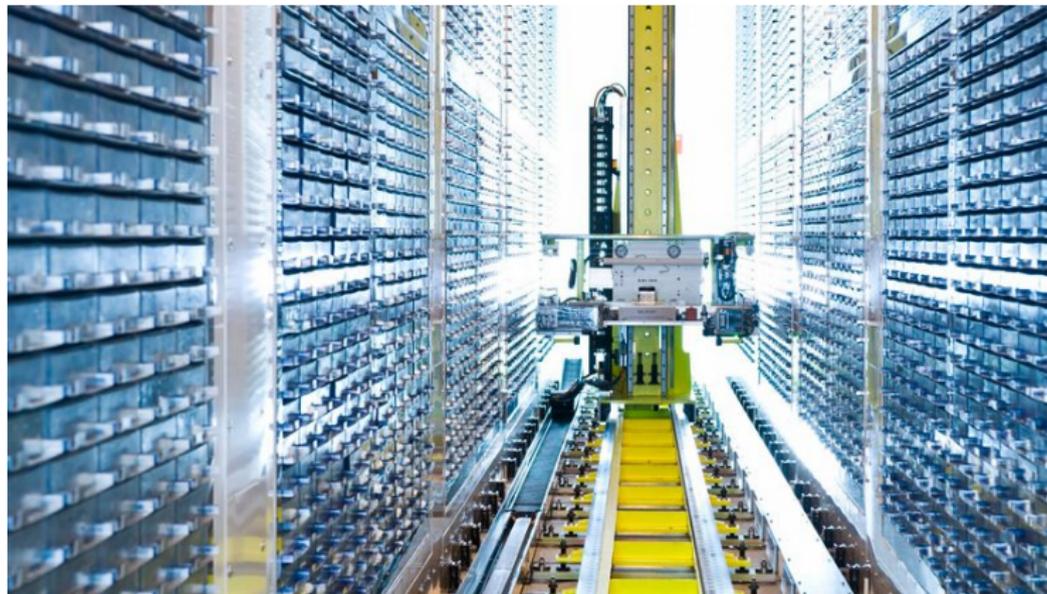
(Image source: Nature)

UK Biobank data

Extensive data on 500,000 individuals, including

- Genotypes
- Diseases (from electronic health records)
- Blood pressure and other clinical diagnostics
- Socioeconomic variables
- Environmental risk factors
- Imaging data
- Diet and exercise questionnaires
- ...

Genotype data

A genotype is an individual's allele at a given *single nucleotide polymorphism* (SNP).

Genotypes measured at 1,000,000 SNPs.
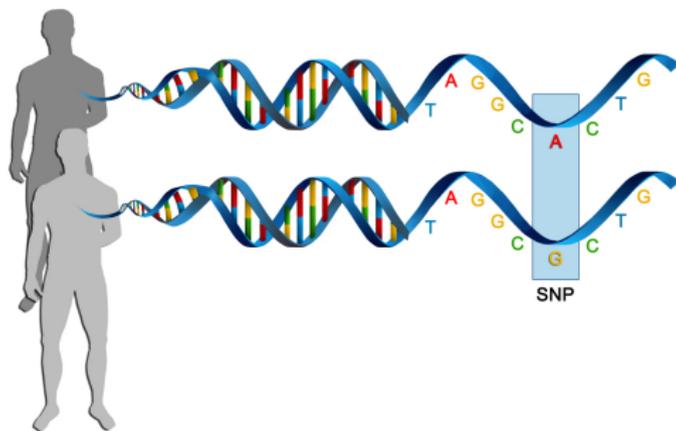
(Image source: Google)

Genotype data have spatial structure

Nearby SNPs are strongly correlated with each other.
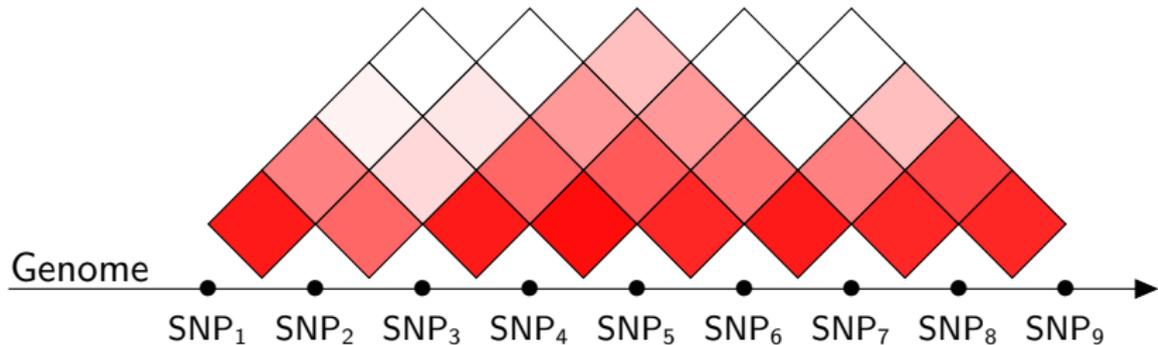
Disease data

Disease codes from hospital episodes,
using *International Classification of Diseases* (ICD-10).

ICD-10 is very comprehensive and includes 20K codes.

Disease data

Disease codes from hospital episodes,
using *International Classification of Diseases* (ICD-10).

ICD-10 is very comprehensive and includes 20K codes.
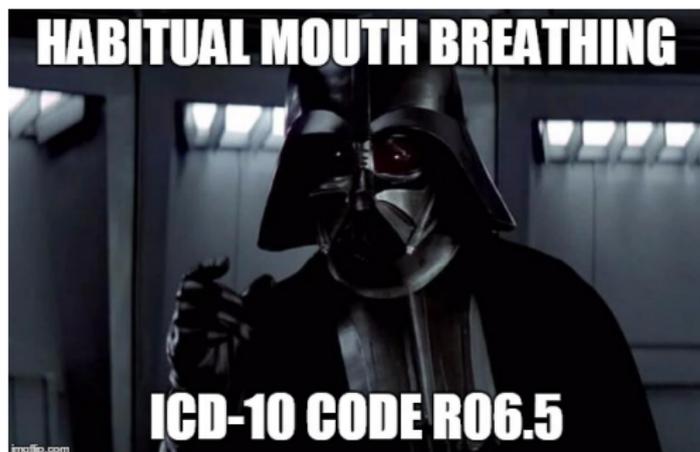
(Image source: Google)

Disease data have tree structure
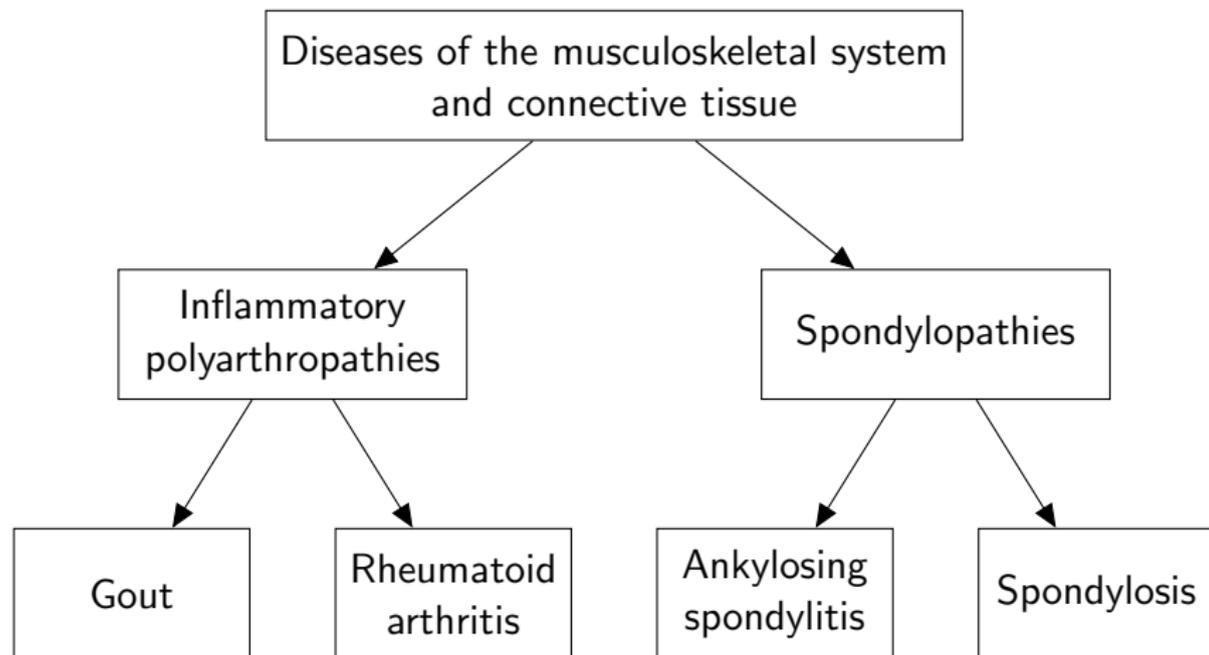
UK Biobank: a complex multiple testing problem
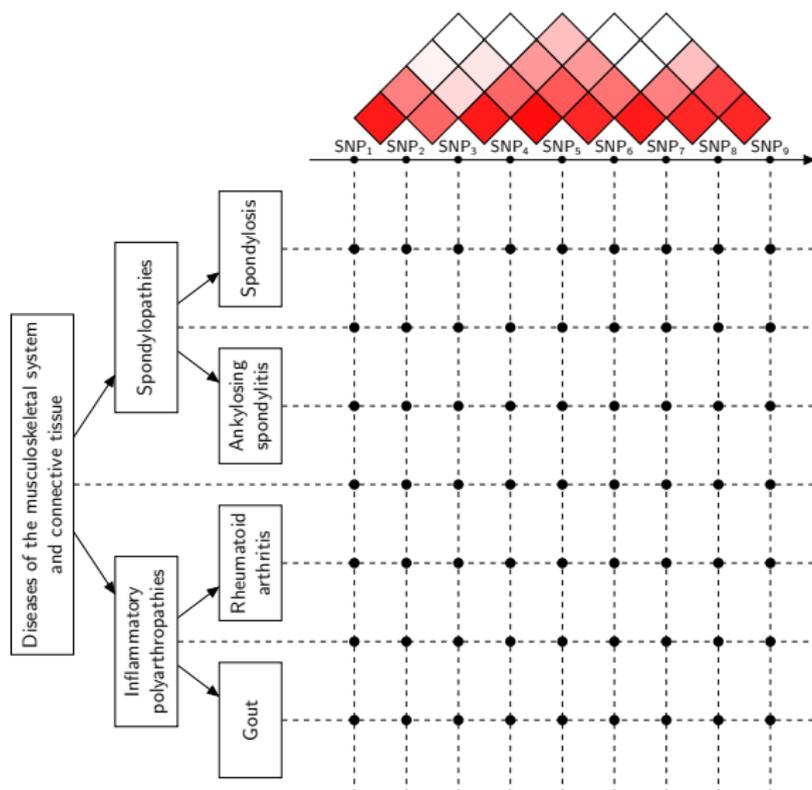
Type-I error rates like the false discovery rate (FDR)
controlled for replicability.

Findings from modern data sets often need curation

Manual curation (exploration):
Domain experts search for interesting patterns in the data.

Automatic curation (filtering):
Structured hypotheses often lead to redundant findings;
filtering is commonly used to reduce redundancy.

Findings from modern data sets often need curation

Manual curation (exploration):
Domain experts search for interesting patterns in the data.

Automatic curation (filtering):
Structured hypotheses often lead to redundant findings;
filtering is commonly used to reduce redundancy.

Curation may conflict with replicability!

Phenome-wide association studies (PheWAS)
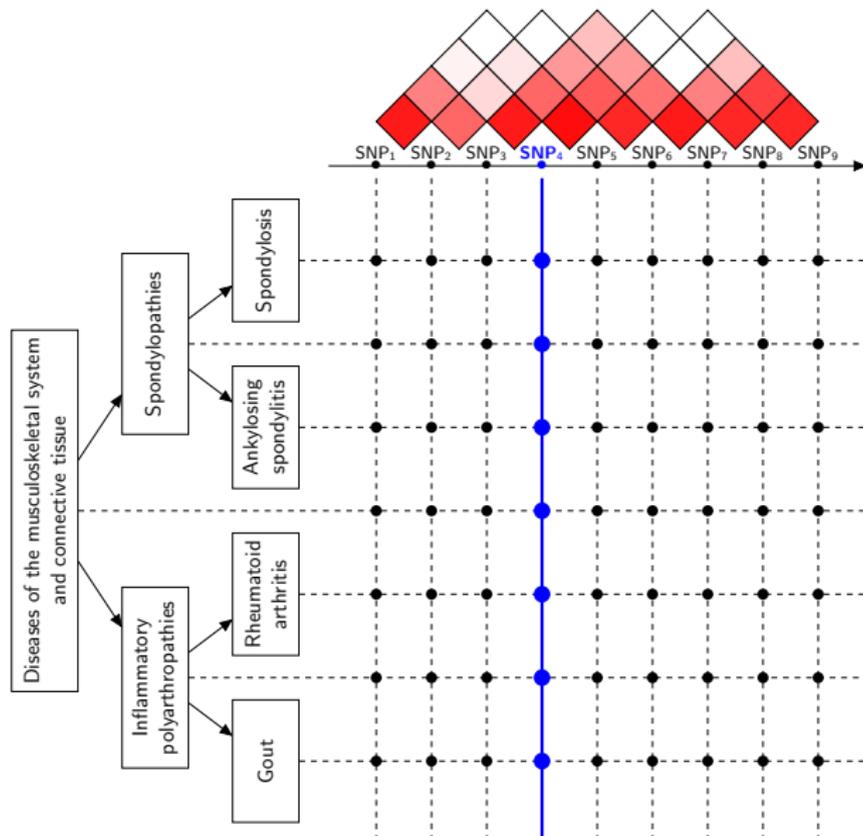
Rejection sets in phenotype space can be redundant
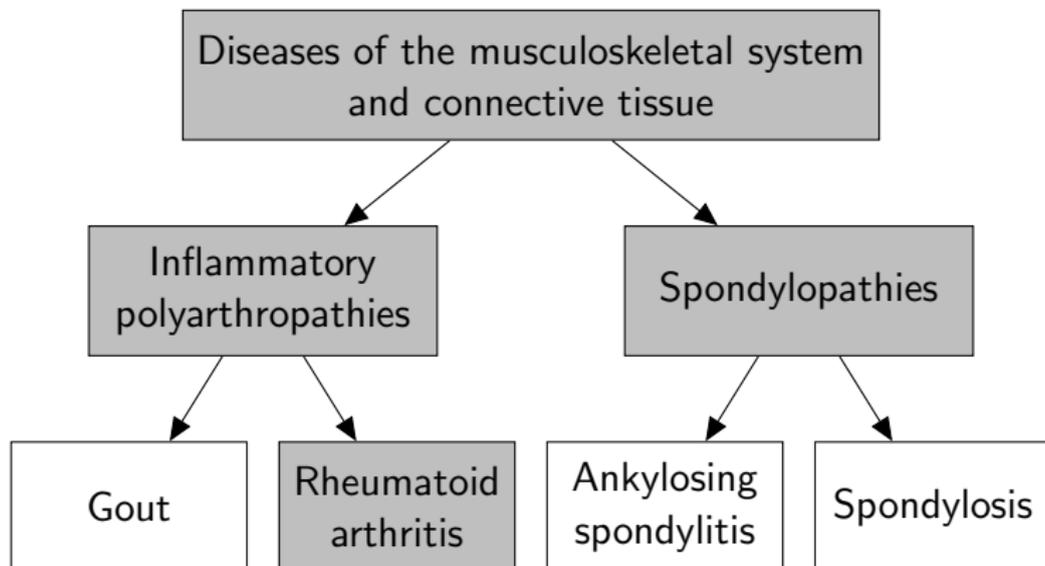
Redundancy can be fixed by applying the *outer nodes filter*
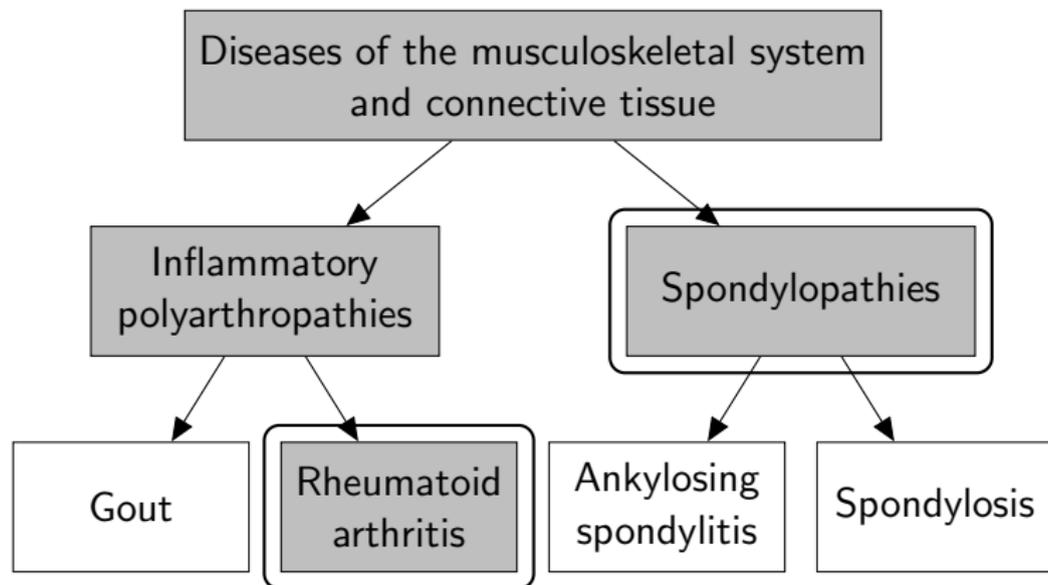
Yekutieli (JASA, 2008)

Outer nodes filter may inflate the FDR
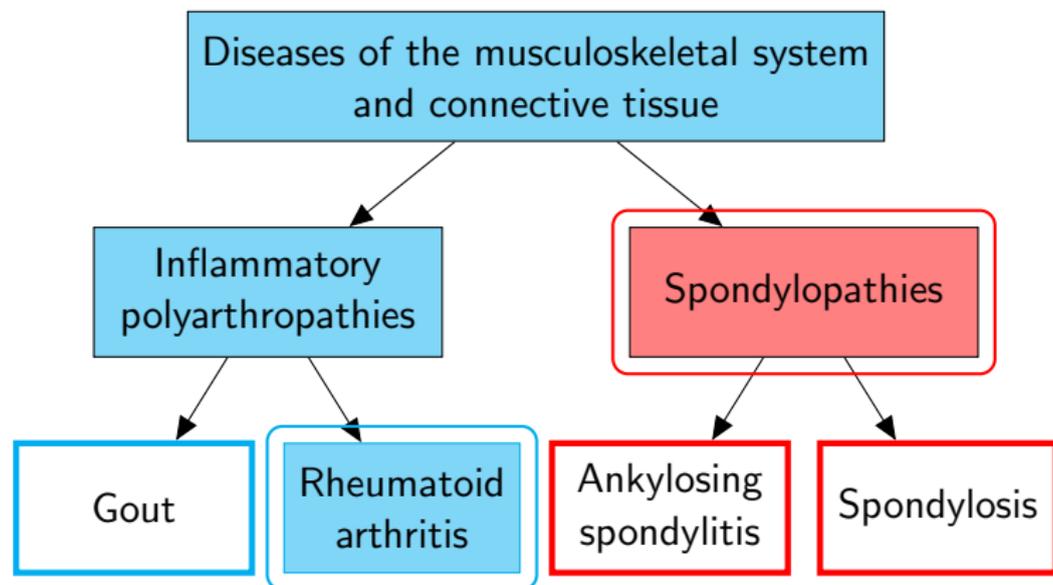
cyan nodes: non-null; red nodes: null; shaded nodes: rejected.

Yekutieli (2008)

Existing options to control outer nodes FDR are limited

- Yekutieli proposed a procedure and bounded its outer nodes FDR, but only under independence.
- Structured Holm procedure[1] controls FWER on DAGs. It allows arbitrary dependence but is conservative.

---

[1]Meijer and Goeman (2016)

Similar problems arise in other applications as well

- ▶ Genome-wide association studies[2]
- ▶ Imaging applications such as fMRI[3]
- ▶ Gene Ontology enrichment analysis[4]
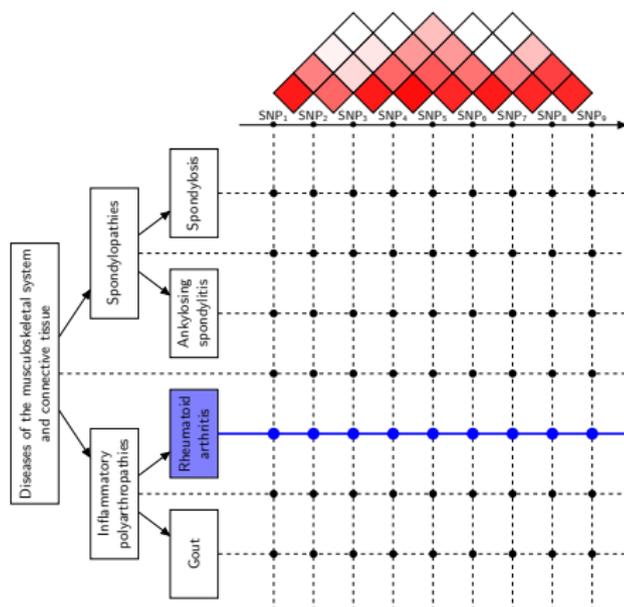
[2]Siegmund, Zhang, Yakir (2011)
[3]Pacifico et al (2004), Heller et al (2006), Sun et al (2015)
[4]Goeman and Bühlmann (2007), Meijer and Goeman (2016)

A general problem

Filtering may inflate the FDR, and must be accounted for.
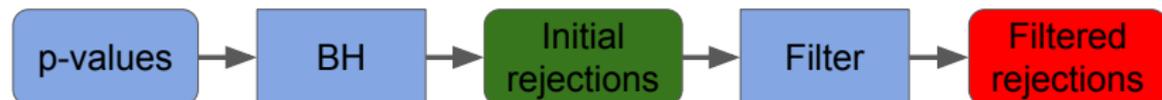
Partial solutions exist, but a general-purpose solution is lacking.

Focus of this talk

Reconciling curation with replicability for
modern data analysis pipelines.

Goeman and Solari (2011), Berk et al (2013), Taylor and
Tibshirani (2015), . . .

Preview: Reconciling curation with replicability

Part I (automatic curation): For any pre-specified filter, we propose **Focused BH**[5] to control the FDR *after filtering*.

---
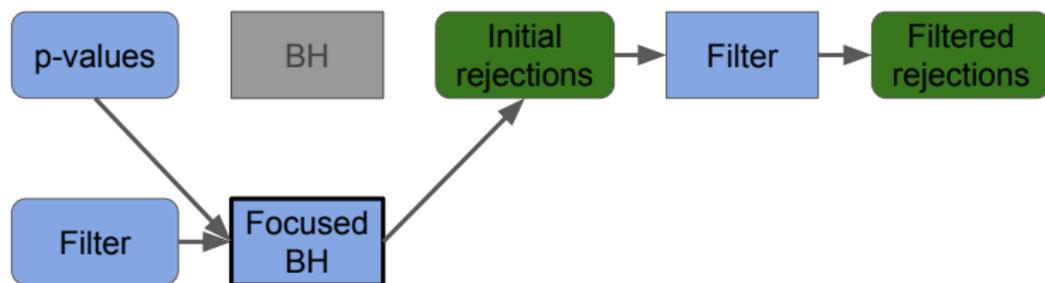
[5]**K.**, Sabatti, Bogomolov (arXiv, 2019+)
[6]**K.** and Ramdas (AOS, in revision, 2019+), **K.** and Sabatti (AOAS, 2019)

Preview: Reconciling curation with replicability

Part I (automatic curation): For any pre-specified filter, we propose **Focused BH**[5] to control the FDR *after filtering*.
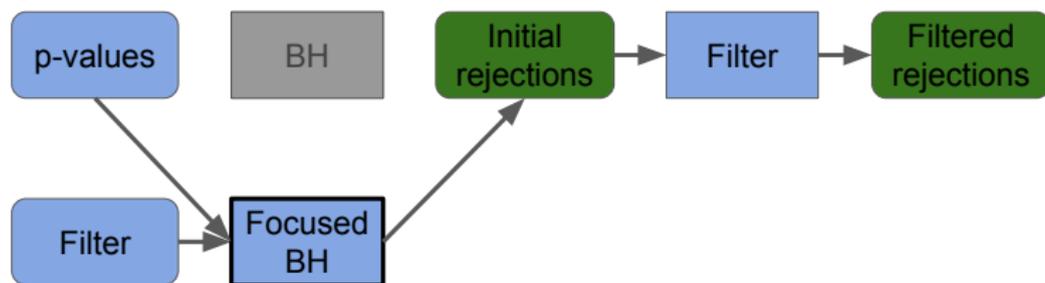
Part II (manual curation): We propose **simultaneous selective inference**[6] to allow directed exploration while bounding FDP whp.

---
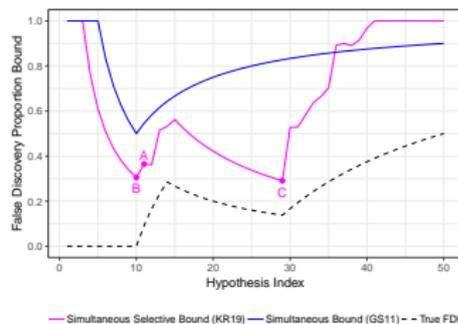
[5]**K.**, Sabatti, Bogomolov (arXiv, 2019+)

[6]**K.** and Ramdas (AOS, in revision, 2019+), **K.** and Sabatti (AOAS, 2019)

Part I: Controlling FDR while filtering

A general definition of a filter

Hypotheses $\mathcal{H} = (H_1, \ldots, H_m)$ and p-values $\boldsymbol{p} = (p_1, \ldots, p_m)$.

Definition

Given $\mathcal{R} \subseteq \mathcal{H}$ and $\boldsymbol{p} \in [0,1]^m$, a *filter* $\mathfrak{F}$ is any mapping

$$\mathfrak{F} : (\mathcal{R}, \boldsymbol{p}) \mapsto \mathcal{U}, \text{ such that } \mathcal{U} \subseteq \mathcal{R}.$$

For example,

- $\mathfrak{F}$ is the outer nodes filter;
- $\mathcal{R}$ is the set of rejected nodes;
- $\mathcal{U}$ is the set of outer nodes.
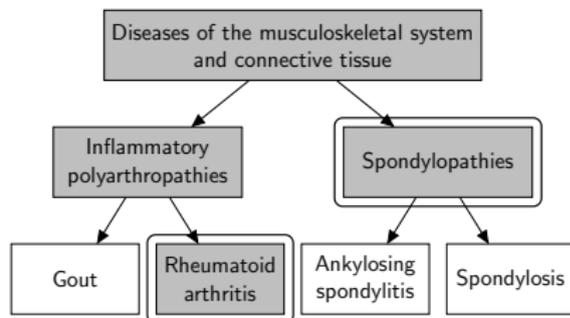
Adjusting the FDR for filtering

The *false discovery proportion* (FDP) of a set $\mathcal{U} \subseteq \mathcal{H}$ is

$$\text{FDP}(\mathcal{U}) = \frac{|\mathcal{U} \cap \mathcal{H}_0|}{|\mathcal{U}|},$$

where $\mathcal{H}_0 \subseteq \mathcal{H}$ is the set of nulls.

Definition

Given a filter $\mathfrak{F}$, the *false filtered discovery rate* of a testing procedure (mapping $\boldsymbol{p} \mapsto \mathcal{R}^*$) is

$$\text{FDR}_{\mathfrak{F}} = \mathbb{E}[\text{FDP}(\mathcal{U}^*)] = \mathbb{E}[\text{FDP}(\mathfrak{F}(\mathcal{R}^*, \boldsymbol{p}))].$$

Given a filter $\mathfrak{F}$ and a pre-specified target FDR level q, our goal is to design a testing procedure for which $\text{FDR}_{\mathfrak{F}} \leq q$.

Adjusting BH to account for filtering

For a p-value cutoff $t \in [0, 1]$, consider $\mathcal{R}(t) = \{j : p_j \leq t\}$.

<div style="background:#eef">

BH procedure

BH employs the FDP estimate (Storey, 2002)

$$\widehat{\mathrm{FDP}}_{\mathsf{BH}}(t) = \frac{m \cdot t}{|\mathcal{R}(t)|};$$

choosing the threshold

$$t^*_{\mathsf{BH}} = \max\{t \in [0, 1] : \widehat{\mathrm{FDP}}_{\mathsf{BH}}(t) \leq q\}.$$

</div>

We are interested instead in $\mathcal{U}(t) = \mathfrak{F}(\{j : p_j \leq t\}, \boldsymbol{p})$.

BH too optimistic in counting discoveries: $|\mathcal{R}(t)| \gg |\mathcal{U}(t)|$.

Adjusting BH to account for filtering

Instead of

$$\widehat{\text{FDP}}_{\text{BH}}(t) = \frac{m \cdot t}{|\mathcal{R}(t)|},$$

correct the denominator and define

$$\widehat{\text{FDP}}(t) = \frac{m \cdot t}{|\mathcal{U}(t)|} = \frac{m \cdot t}{|\mathfrak{F}(\{j : p_j \leq t\}, \boldsymbol{p})|}.$$

We keep the numerator as is, since $|\mathcal{U}(t) \cap \mathcal{H}_0| \leq |\mathcal{R}(t) \cap \mathcal{H}_0|$.

Focused BH procedure

Data: p-values $p_1, \ldots, p_m$, filter $\mathfrak{F}$, target level q

for $t \in \{0, p_1, \ldots, p_m\}$ **do**

    Compute $\widehat{\mathrm{FDP}}(t) = \dfrac{m \cdot t}{|\mathfrak{F}(\{j : p_j \leq t\}, \boldsymbol{p})|}$;

end

Compute $t^* \equiv \max\{t \in \{0, p_1, \ldots, p_m\} : \widehat{\mathrm{FDP}}(t) \leq q\}$;

Result: $\mathcal{R}^* = \{j : p_j \leq t^*\}$.

- Focused BH is a general-purpose way of dealing with filters; note that $\mathfrak{F}$ can be a black box.
- When $\mathfrak{F}$ does nothing, Focused BH reduces to BH.
- Procedure can be expanded to filters that *prioritize* rejections.

Focused BH provably controls FDR$_\mathfrak{F}$

A filter $\mathfrak{F}$ is **monotonic** if for $\mathcal{R}^1 \supseteq \mathcal{R}^2$ and $\boldsymbol{p}^1 \leq \boldsymbol{p}^2$, we have

$$|\mathfrak{F}(\mathcal{R}^1, \boldsymbol{p}^1)| \geq |\mathfrak{F}(\mathcal{R}^2, \boldsymbol{p}^2)|.$$

A filter is **simple** if $|\mathfrak{F}(\mathcal{R}, \boldsymbol{p})|$ is independent of $\boldsymbol{p}$.

> ### Theorem (**K.**, Sabatti, Bogomolov)
>
> Focused BH controls FDR$_\mathfrak{F}$ if either
> 1. p-values are independent, $\mathfrak{F}$ is simple or monotonic.
> 2. p-values are "positively dependent" (PRDS), $\mathfrak{F}$ is monotonic.

- ▶ Proof for item 1 inspired by Benjamini and Bogomolov (2014);
- ▶ Proof for item 2 inspired by Blanchard and Roquain (2008).

Simulations suggest Focused BH is robust.

Specializing to the outer nodes filter

Corollary

Focused BH controls the outer nodes FDR on trees if
the p-values are positively dependent.

Proof: The outer nodes filter is monotonic on trees.
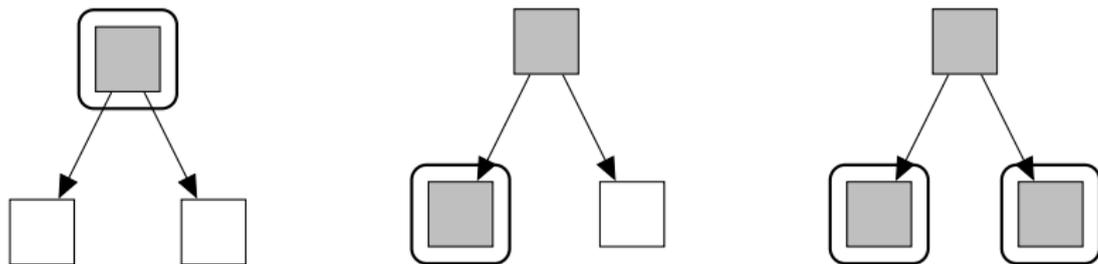
Specializing to the outer nodes filter

Corollary

Focused BH controls the outer nodes FDR on trees if the p-values are positively dependent.

Proof: The outer nodes filter is monotonic on trees.
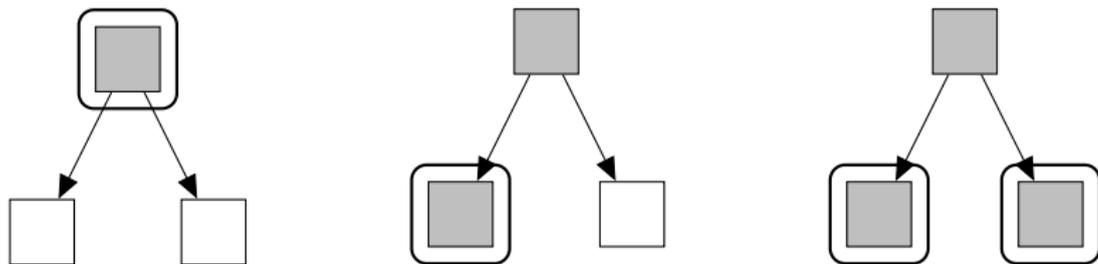
Focused BH is the first procedure provably controlling outer nodes FDR under dependence.

Improving the power of Focused BH

The numerator $m \cdot t$ in

$$\widehat{\mathrm{FDP}}(t) = \frac{m \cdot t}{|\mathfrak{F}(\{j : p_j \leq t\}, \boldsymbol{p})|}$$

can be a conservative estimate of $V(t) = |\mathcal{U}(t) \cap \mathcal{H}_0|$.

Can improve procedure's power by tightening FDP estimate, e.g.

$$\widehat{V}_{\mathrm{oracle}}(t) = \mathbb{E}[V(t)] \leq m \cdot t.$$

Improving the power of Focused BH by permutations

Let $\widetilde{\boldsymbol{p}}$ be a "permuted" version of $\boldsymbol{p}$. Then,

$$\mathbb{E}[V(t)] = \mathbb{E}\left[|\mathfrak{F}(\{j : p_j \leq t\}, \boldsymbol{p}) \cap \mathcal{H}_0|\right]$$
$$\approx \mathbb{E}[|\mathfrak{F}(\{j : \widetilde{p}_j \leq t\}, \widetilde{\boldsymbol{p}}) \cap \mathcal{H}_0|]$$
$$\leq \mathbb{E}[|\mathfrak{F}(\{j : \widetilde{p}_j \leq t\}, \widetilde{\boldsymbol{p}})|].$$

Given permutations $\widetilde{\boldsymbol{p}}^1, \ldots, \widetilde{\boldsymbol{p}}^B$, define

$$\widehat{V}_{\mathsf{perm}}(t) = \frac{1}{B} \sum_{b=1}^{B} |\mathfrak{F}(\{j : \widetilde{p}_j^b \leq t\}, \widetilde{\boldsymbol{p}}^b)|.$$

No theoretical results yet, but performs well in simulations.

Simulation: Setup

Graph structure: Forest of 20 binary trees of depth 6, with $m = 1260$ total nodes.

Data generating mechanism:

- 21 non-null leaves (out of 640), 98 total non-nulls;
- Leaf nodes get independent p-values;
- Internal nodes get p-values by applying Simes global test to their leaf descendants.

Filter: Outer nodes filter.

Simulation: Methods compared

- BH   (targeting pre-filter FDR at level $q = 0.1$)
- Structured Holm[7] (targeting FWER at level $q = 0.1$)
- Yekutieli[8]   (targeting post-filter FDR at level $q = 0.1$)
- Focused BH   (targeting post-filter FDR at level $q = 0.1$)
  - Original version
  - Permutation version
  - Oracle version

---

[7] Meijer and Goeman (2016)

[8] Yekutieli (2008)

Simulation: Results

Application: UK Biobank PheWAS with outer nodes filter

HLA region on chromosome 6 is known to affect many diseases.

Conducted PheWAS analysis for the HLA-B*27:05 allele, studied previously by Cortes et al (Nature Genetics, 2017).

Computed p-values testing marginal association between this allele and the $m = 3265$ ICD-10 codes that had at least 50 cases.[9]
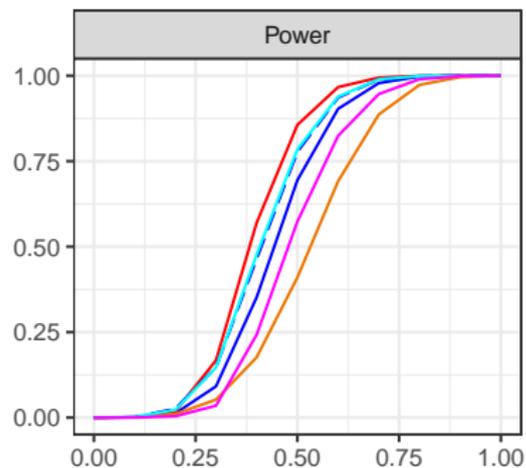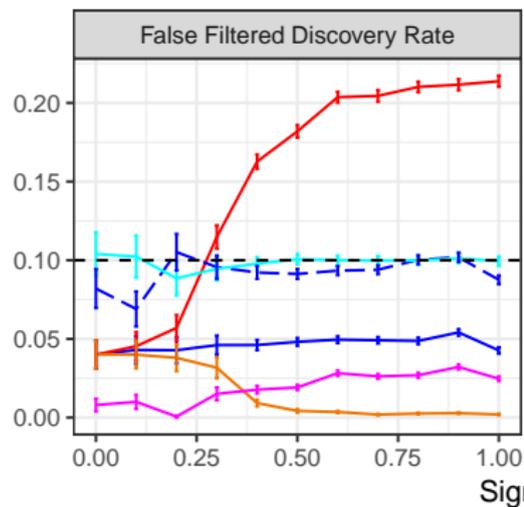
BH, Structured Holm, Yekutieli, Focused BH applied with $q = 0.05$.

---

[9]This filtering step does not need to be corrected for, since it does not take the response variable into account.

Number of outer node rejections made by each method

Method	Outer node rejections
BH	28
Focused BH	**24**
Structured Holm	13
Yekutieli	1

Focused BH rejects 34 nodes, 24 outer nodes

Key
- Not rejected by Focused BH
- Rejected by Focused BH
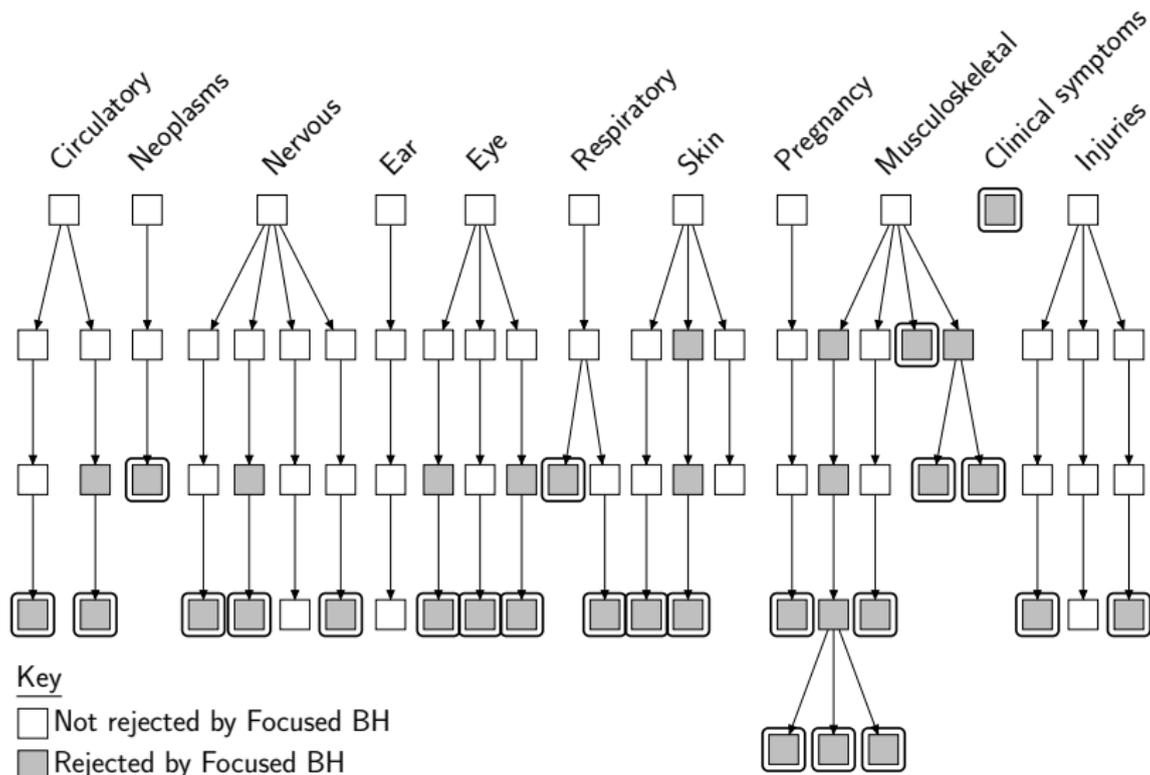
FBH rejects 11 outer nodes more than Structured Holm

Key
- ☐ Not rejected by Focused BH
- ▨ Rejected by Focused BH
- ◼ Outer node rejected by Focused BH but not Structured Holm
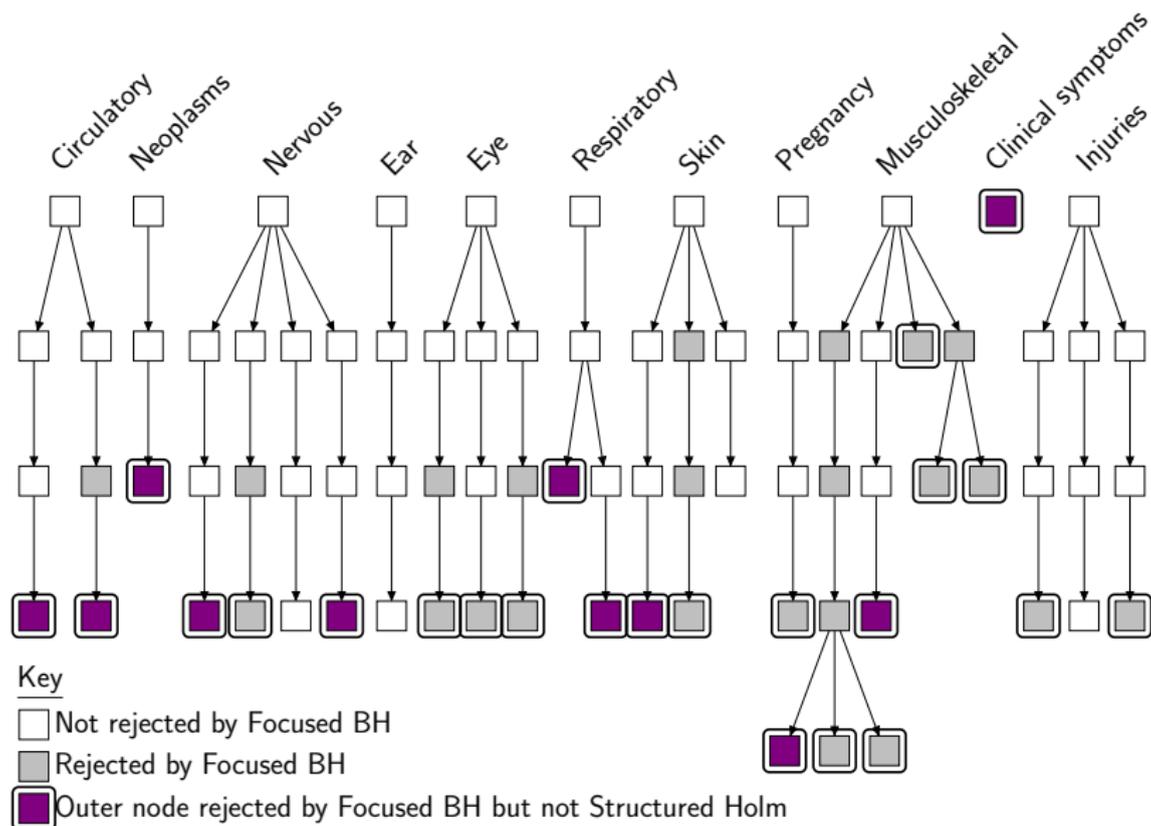
Summary of Focused BH
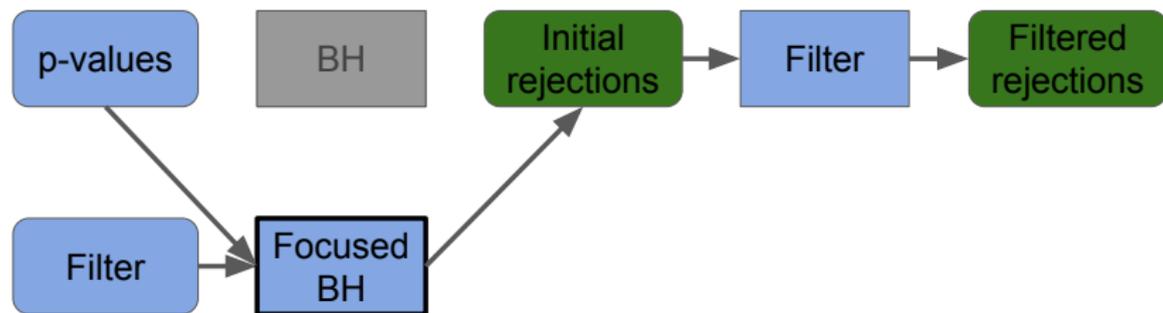
Focused BH guarantees Type-I error control when data analysis involves automatic curation via a pre-specified filter.

Filtering framework is general; applies beyond examples presented.

Part II: From automatic to manual curation

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \ldots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \ldots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

🔍

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \dots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \ldots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \ldots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

Manually curating promising hypotheses

Consider the practice of re-running an FDR procedure with different target levels until one obtains a "good" rejection set.

$\mathcal{R}_k = \{H_{(1)}, \ldots, H_{(k)}\}$: set corresponding to k smallest p-values.

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}$$

Simultaneous inference is one solution (e.g. Goeman and Solari 2011, Berk et al 2013), but can be conservative.

Simultaneous selective inference

Data scientist wants to inspect a "menu" of options

$$\varnothing = \mathcal{R}_0 \subseteq \mathcal{R}_1 \subseteq \cdots \subseteq \mathcal{R}_m \subseteq \mathcal{H}.$$

Idea: provide corresponding upper bounds

$$\overline{\mathrm{FDP}}(\mathcal{R}_k) = \frac{\log(\alpha^{-1})}{\log(1 + \log(\alpha^{-1}))} \frac{1 + n \cdot p_{(k)}}{|\mathcal{R}_k|}$$

such that


Theorem (K. and Ramdas, AOS, in revision, 2019+)

Under independence of null p-values,

$$\mathbb{P}[\mathrm{FDP}(\mathcal{R}_k) \leq \overline{\mathrm{FDP}}(\mathcal{R}_k) \text{ for all } k] \geq 1 - \alpha$$

for all n and all $\alpha \leq 0.31$.

Data scientist can freely choose from menu while maintaining validity of FDP bounds.

Simultaneous selective inference in a toy example
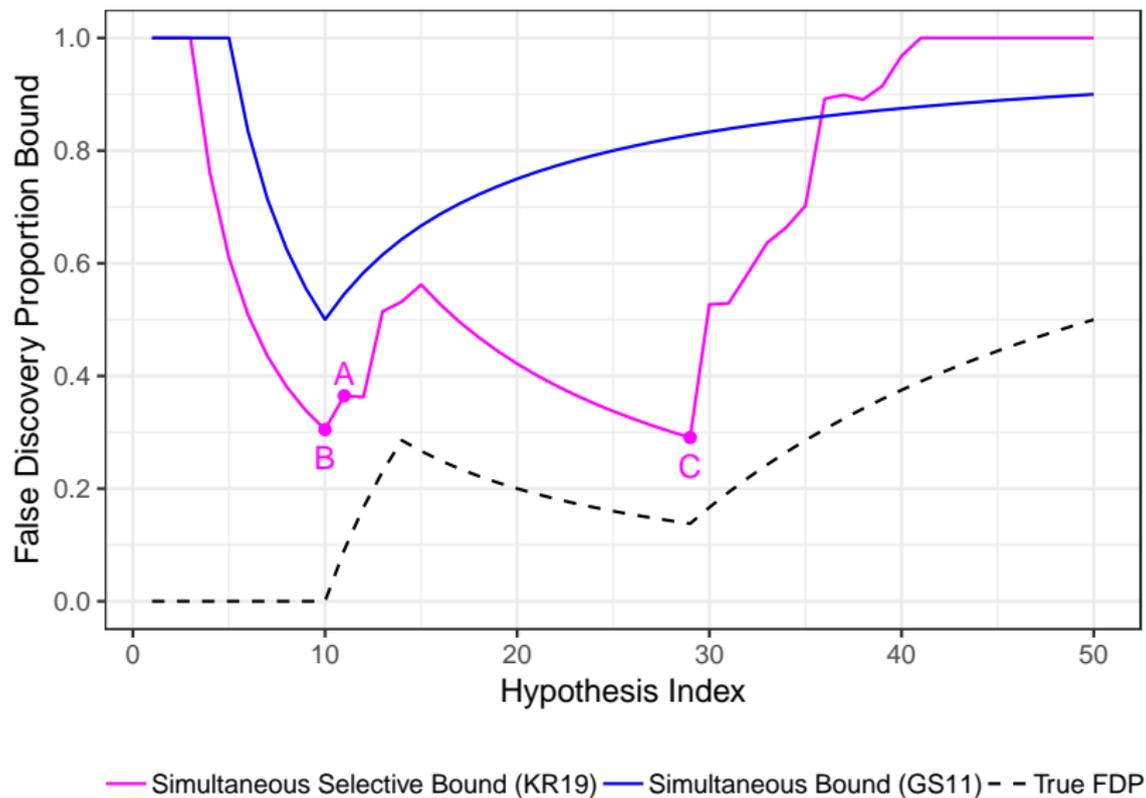
Linear upper bounds for empirical processes

For bounds of the form $\overline{\mathrm{FDP}}(t) = \dfrac{a + bt}{R(t)}$, we seek a, b such that

$$\mathbb{P}[V(t) \leq a + bt \text{ for all } t \in [0, 1]] \geq 1 - \alpha,$$

where $V(t) = \sum_{j \in \mathcal{H}_0} I(p_j \leq t)$.

[10]Robbins (1954) and Dvoretsky Kiefer Wolfowitz (1956)

Linear upper bounds for empirical processes

For bounds of the form $\overline{\text{FDP}}(t) = \dfrac{a + bt}{R(t)}$, we seek a, b such that

$$\mathbb{P}[V(t) \leq a + bt \text{ for all } t \in [0, 1]] \geq 1 - \alpha,$$

where $V(t) = \sum_{j \in \mathcal{H}_0} I(p_j \leq t)$.

Existing finite-sample bounds:[10]

- $\overline{V}(t) = \frac{1}{\alpha} nt$;
  tight very near 0.

- $\overline{V}(t) = \sqrt{\frac{n}{2} \log \frac{1}{\alpha}} + nt$;
  tight near 1.

We obtain a new bound by
exploiting connection between
empirical and Poisson processes.

Empirical process hitting linear boundary

---
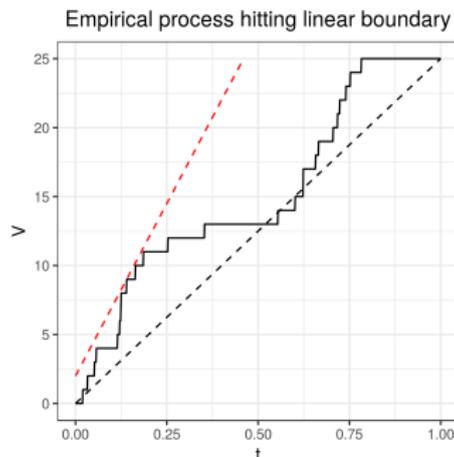
[10] Robbins (1954) and Dvoretsky Kiefer Wolfowitz (1956)

Comparing to existing bounds ($n = 500, \alpha = 0.05$)
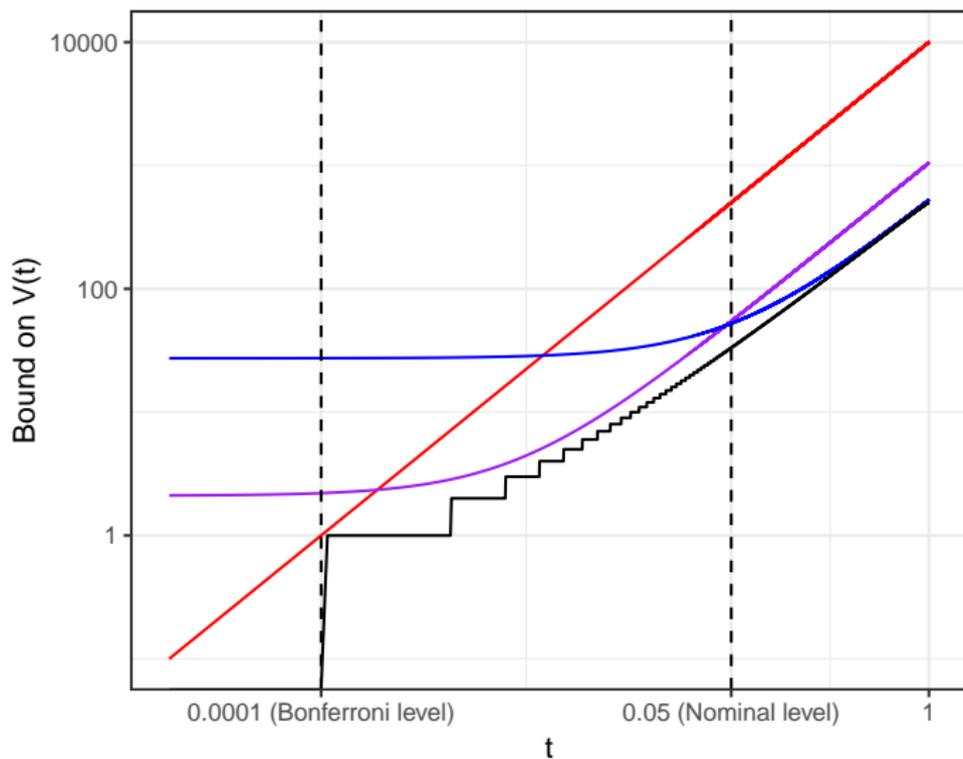
Simultaneous selective inference with side information

KR19+ bounds can leverage side information to give data scientists a better menu of rejection sets to choose from.

- ▶ Hypotheses ordered a priori
  (same menu as accumulation test[11])

- ▶ Hypotheses ordered adaptively
  (same menu as AdaPT or STAR[12])

- ▶ Hypotheses ordered according to variable selection importance
  (same menu as knockoffs[13])

---

[11]Li and Barber (2017)
[12]Lei and Fithian (2018), Lei, Ramdas, Fithian (2019+)
[13]Barber and Candes (2015)

Simultaneous selective inference for knockoffs

Knockoffs method (Barber and Candes, 2015) developed for variable selection with FDR control.

Knockoff statistics $W_1, \ldots, W_m$ assigned to variables instead of p-values, ordering variables based on

$$W_{(1)} \geq W_{(2)} \geq \cdots \geq W_{(m)}.$$

BR19+ derived uniform FDP bounds for knockoffs as well:

$$\overline{\mathrm{FDP}}(\mathcal{R}_k) = \frac{\log(\frac{1}{\alpha})}{\log(2 - \alpha)} \frac{1 + |\{j : W_j \leq -W_{(k)}\}|}{|\mathcal{R}_k|}.$$

Uniform bounds for knockoffs first considered by **K.** and Sabatti (AOAS, 2019).

Replicability guarantees for modern data analysis pipelines

Different modes of curation require different statistical approaches:

Mode of curation	Statistical approach
1. Automatic (filtering)	Focused BH
2. Manual (exploration)	Simultaneous selective inference

Replicability guarantees for modern data analysis pipelines

Different modes of curation require different statistical approaches:

Mode of curation	Statistical approach
1. Automatic (filtering)	Focused BH
2. Manual (exploration)	Simultaneous selective inference

These lie on a spectrum from selective to simultaneous inference:
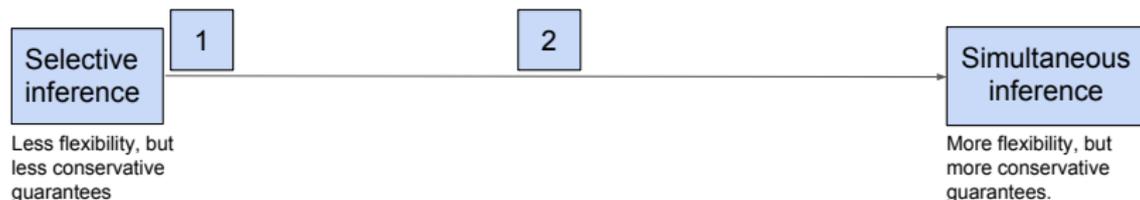
Replicability guarantees for modern data analysis pipelines

Different modes of curation require different statistical approaches:

Mode of curation	Statistical approach
1. Automatic (filtering)	Focused BH
2. Manual (exploration)	Simultaneous selective inference

These lie on a spectrum from selective to simultaneous inference:
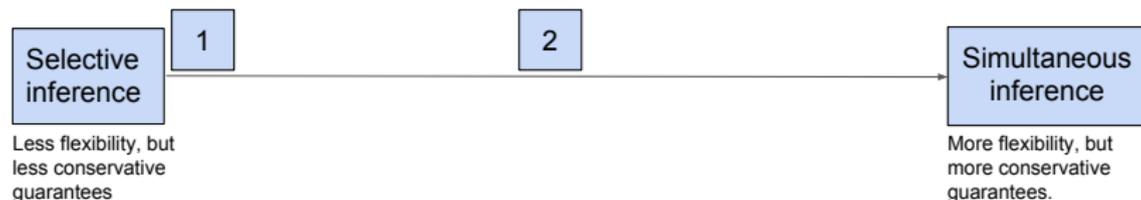
Open questions:

- ▶ (Applications) Pairing applications with inferential guarantees;
- ▶ (Theory, Methodology) Filling in the spectrum with powerful procedures using realistic assumptions.

Thank you.

All papers and code available at
`http://web.stanford.edu/~ekatsevi/index.html`.

PRDS condition

Definition (Benjamini Yekutieli 2001)

The vector $\boldsymbol{p}$ is PRDS if for any null j and non-decreasing set $\mathcal{D} \subseteq [0,1]^m$, the quantity $\mathbb{P}[\boldsymbol{p} \in \mathcal{D} | p_i \leq t]$ is nondecreasing in $t \in (0,1]$.

Definition of power in the context of filtering

Maximum possible weighted number of non-null rejections is

$$T_{\mathsf{max}} \equiv \max_{\mathcal{R}, \boldsymbol{p}} \left\{ \sum_{j \in \mathcal{H}_1} U_j \right\}; \quad \boldsymbol{U} = \mathfrak{F}(\mathcal{R}, \boldsymbol{p}),$$

Then, define power via

$$\pi(\boldsymbol{U}) = \mathbb{E} \left[\frac{\sum_{j \in \mathcal{H}_1} U_j}{T_{\mathsf{max}}} \right].$$

# Simulation 2: GWAS with clump filtering
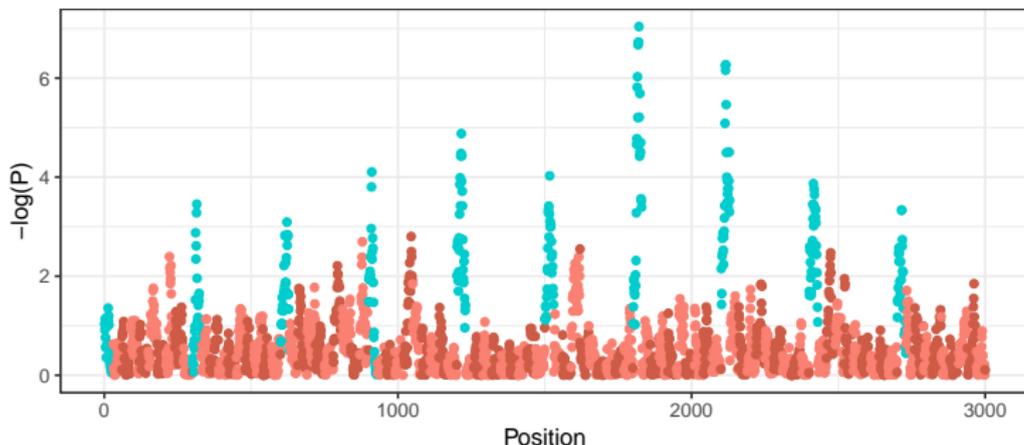
- ▶ Genome of length 3000, with 100 LD blocks of size 30
- ▶ Simulated genotype data with local correlations
- ▶ Phenotypes from linear model with 10 nonzero coefficients
- ▶ Univariate association p-values generated for each SNP
- ▶ For simplicity, filter uses a priori LD blocks as clumps

Simulation 2: Results
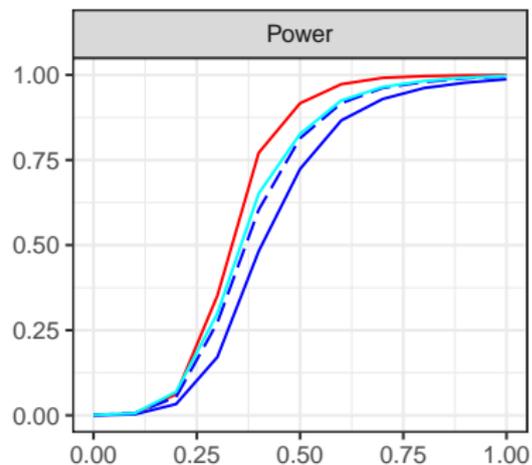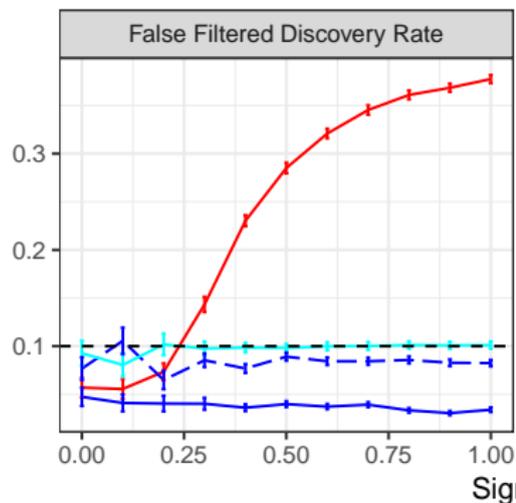
Robustness experiment
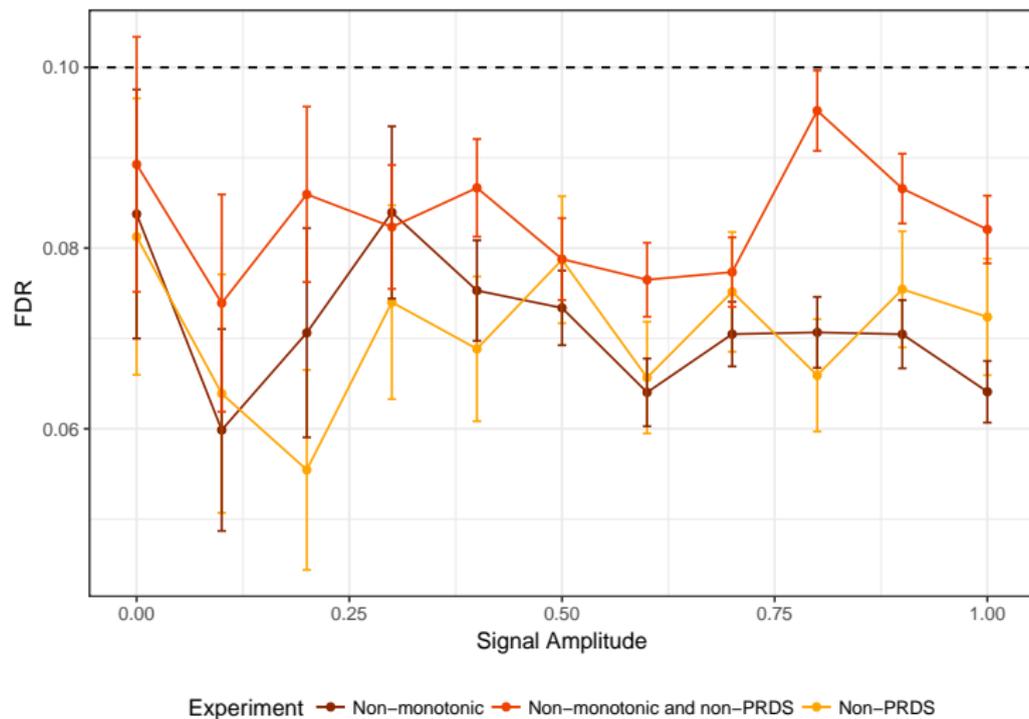
Outer nodes found by BH but not Focused BH

- Other and unspecified antidepressants [as a cause of death via complication of medical care]
- Urticaria [also known as hives]
- Localisation-related (focal) (partial) symptomatic epilepsy and epileptic syndromes with complex partial seizures
- Meniere's disease

Outer nodes found by Focused BH but not Structured Holm

- Symptoms, signs and abnormal clinical and laboratory findings
- Other benign neoplasms of connective and other soft tissues
- Meningitis, unspecified
- Other specified polyneuropathies
- Cardiomegaly
- Scrotal varices
- Chronic sinusitis
- Paralysis of vocal cords and larynx
- Cellulitis of other sites
- Rheumatoid arthritis, unspecified (Multiple sites)
- Other synovitis and tenosynovitis

FBH rejects 4 nodes fewer than BH

Key
- □ Not rejected by Focused BH
- ▨ Rejected by Focused BH
- ▨ Rejected by BH but not Focused BH
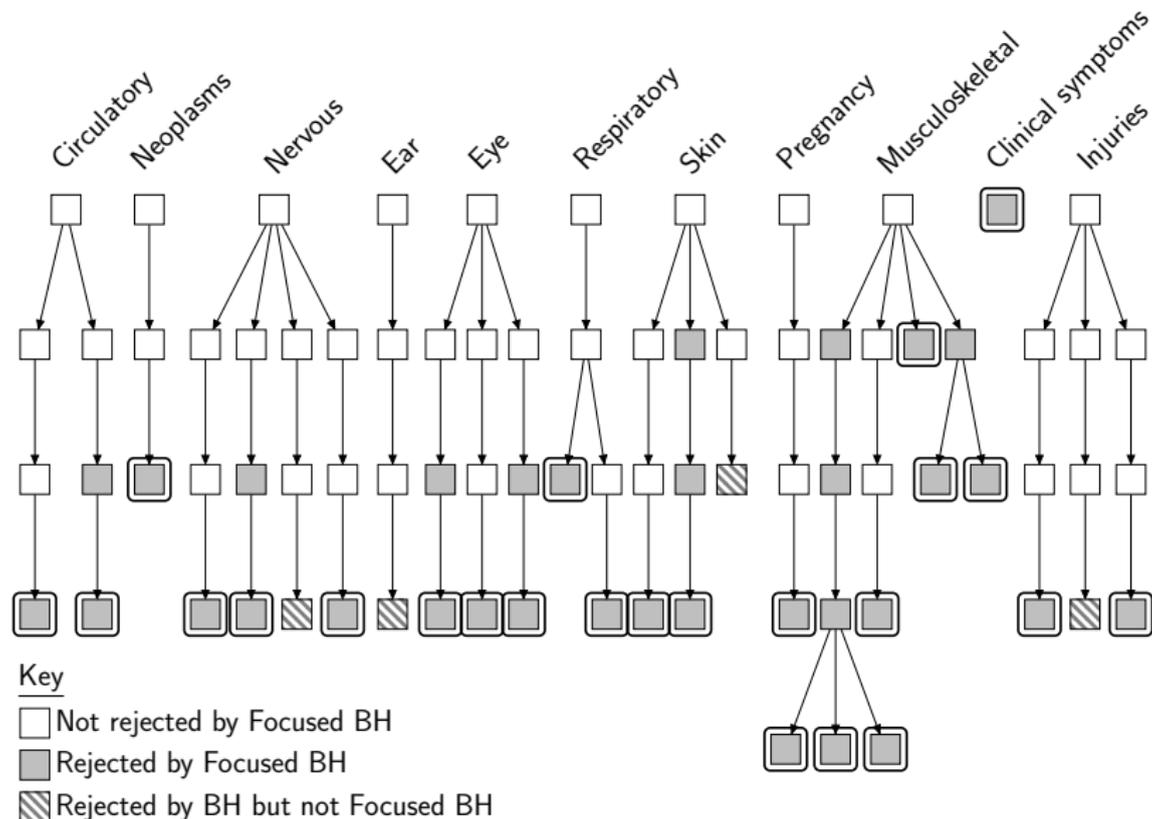
Focusing on diseases of the musculoskeletal system
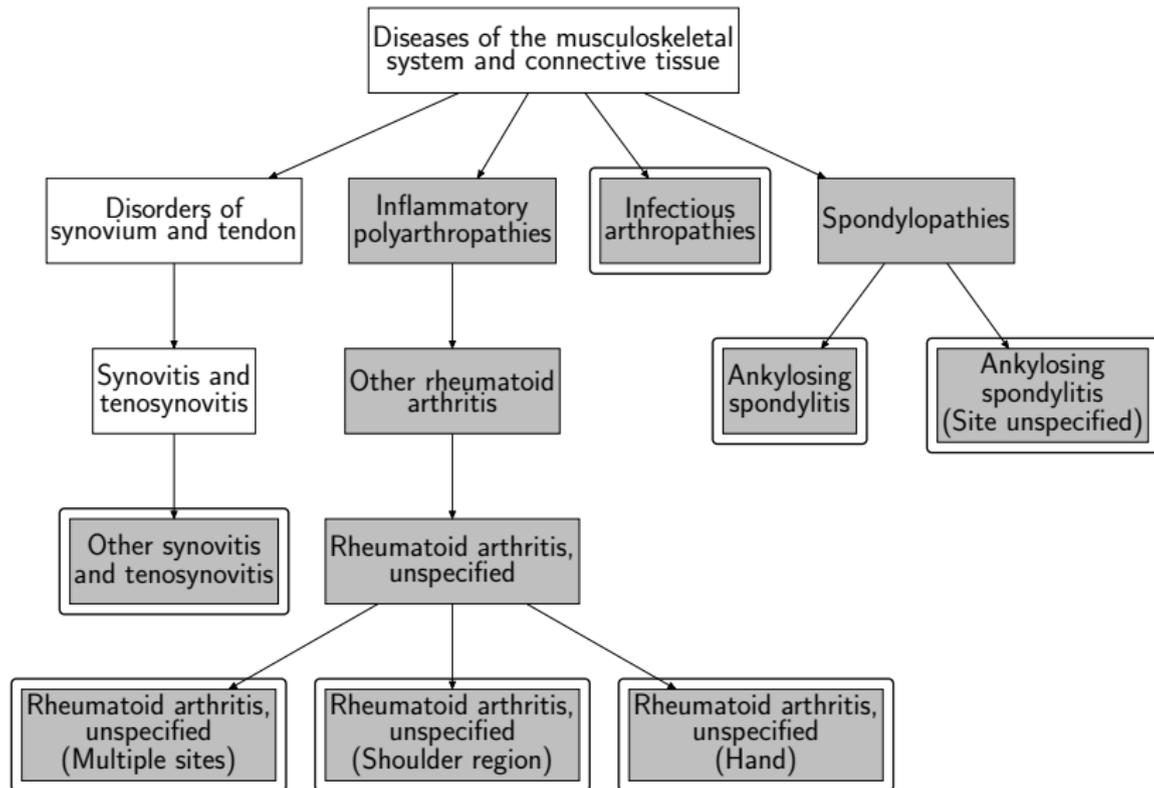
Focusing on diseases of the skin
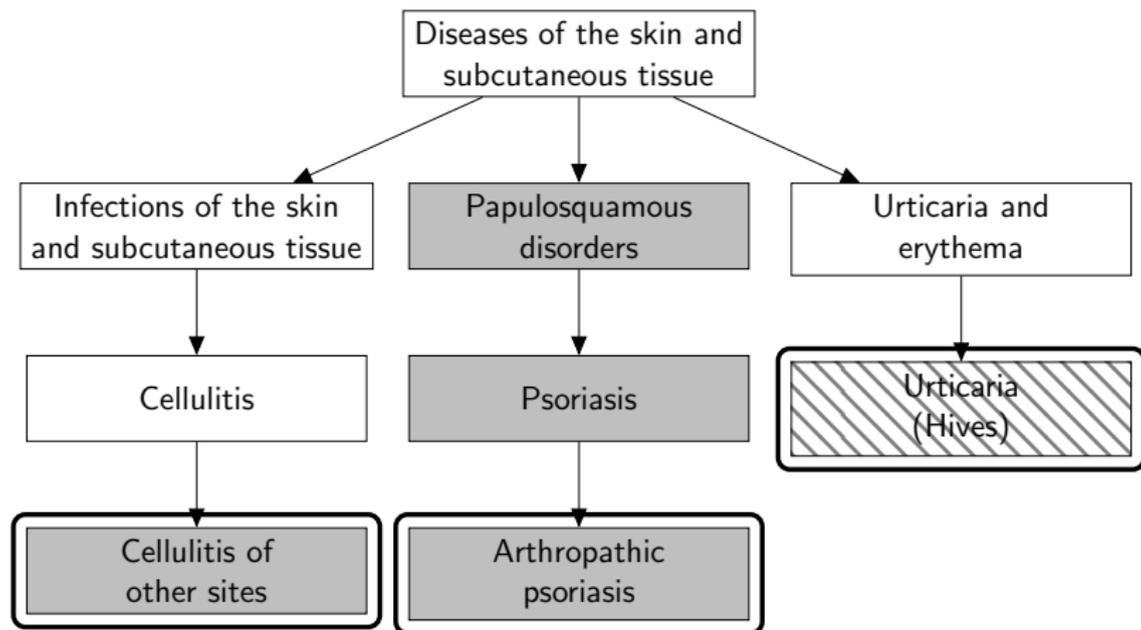
Soft outer nodes filter
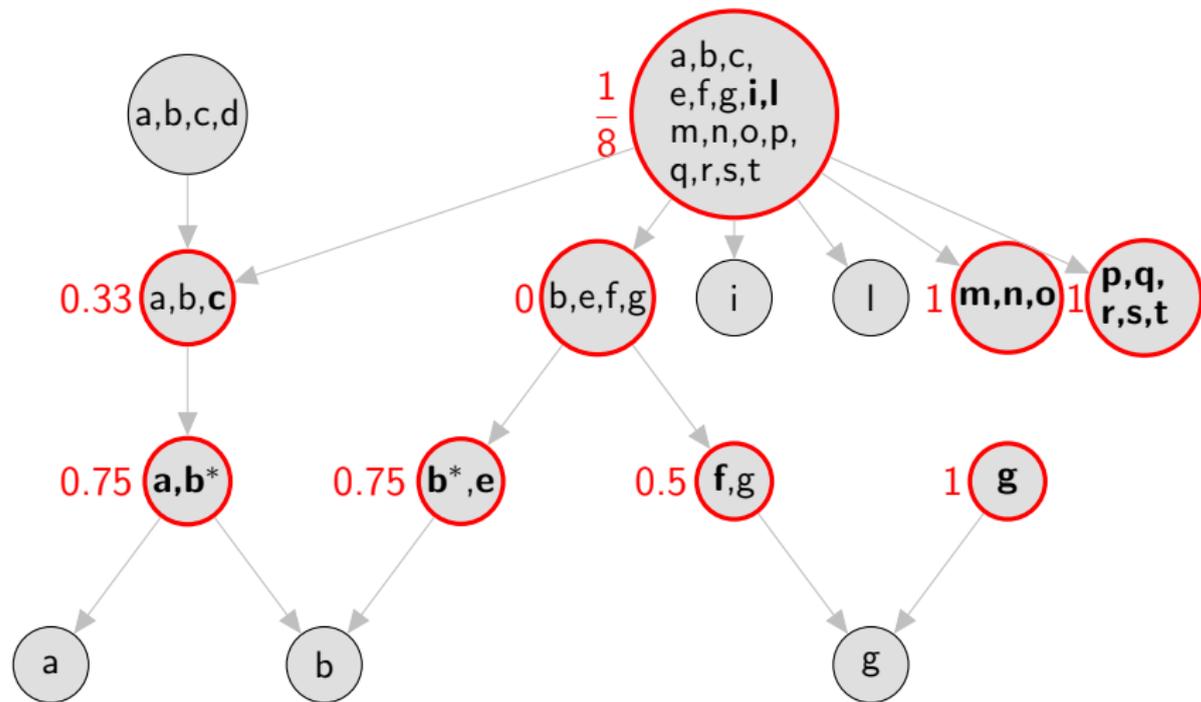
Multi-filter Focused BH

Given M filters $\mathfrak{F}_1, \ldots, \mathfrak{F}_M$, suppose one wants $\mathcal{R}^*$ such that

$$\text{FDP}_{\mathfrak{F}_k} = \mathbb{E}[\text{FDP}(\mathfrak{F}_k(\mathcal{R}^*, \boldsymbol{p}))] \leq q_k \text{ for all } k = 1, \ldots, m.$$

For a threshold t, we can construct $\widehat{\text{FDP}}_k(t)$ as in Focused BH, and then choose

$$t^* = \max\{t \in \{0, p_1, \ldots, p_m\} : \widehat{\text{FDP}}_k(t) \leq q_k \text{ for all } k\}.$$

This will control FDR for all filtered rejection sets if $\boldsymbol{p}$ is PRDS and all filters are monotonic.

Focused Storey BH

Writing

$$\widehat{m}_0^\lambda = \frac{1 + |\{j : p_j > \lambda\}|}{1 - \lambda},$$

following Storey, we can define

$$\widehat{\mathrm{FDP}}_{\mathsf{Storey}}(t) = \frac{\widehat{m}_0^\lambda \cdot t}{|\mathfrak{F}(\mathcal{R}(t, \boldsymbol{p}), \boldsymbol{p})|}.$$

The corresponding procedure controls FDR under independence for simple filters.