Optimizing the Relevance-Redundancy Tradeoff for Efficient Semantic Segmentation

Caner Hazırbaş

Joint work with Julia Diebold and Daniel Cremers

Optimizing the Relevance-Redundancy Tradeoff for Efficient Semantic Segmentation

Caner Hazırbaş

$$\min \mathsf{E} = \min \{\lambda \cdot \boxed{\mathsf{D}} + \mathsf{R}\}$$

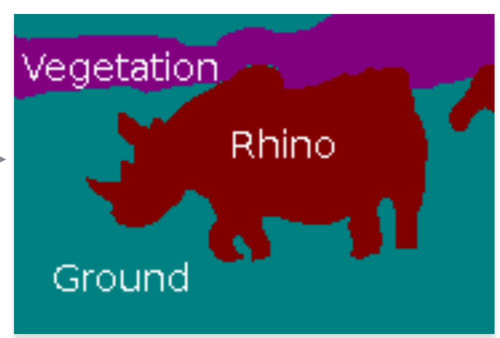

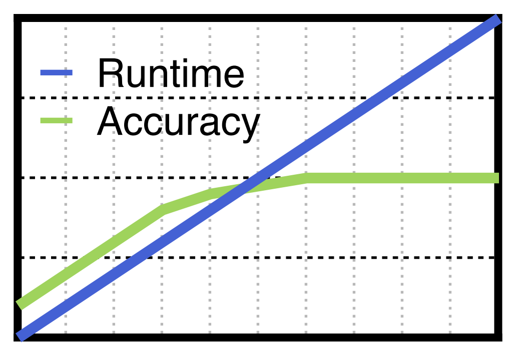

Joint work with Julia Diebold and Daniel Cremers

Optimizing the Relevance-Redundancy Tradeoff

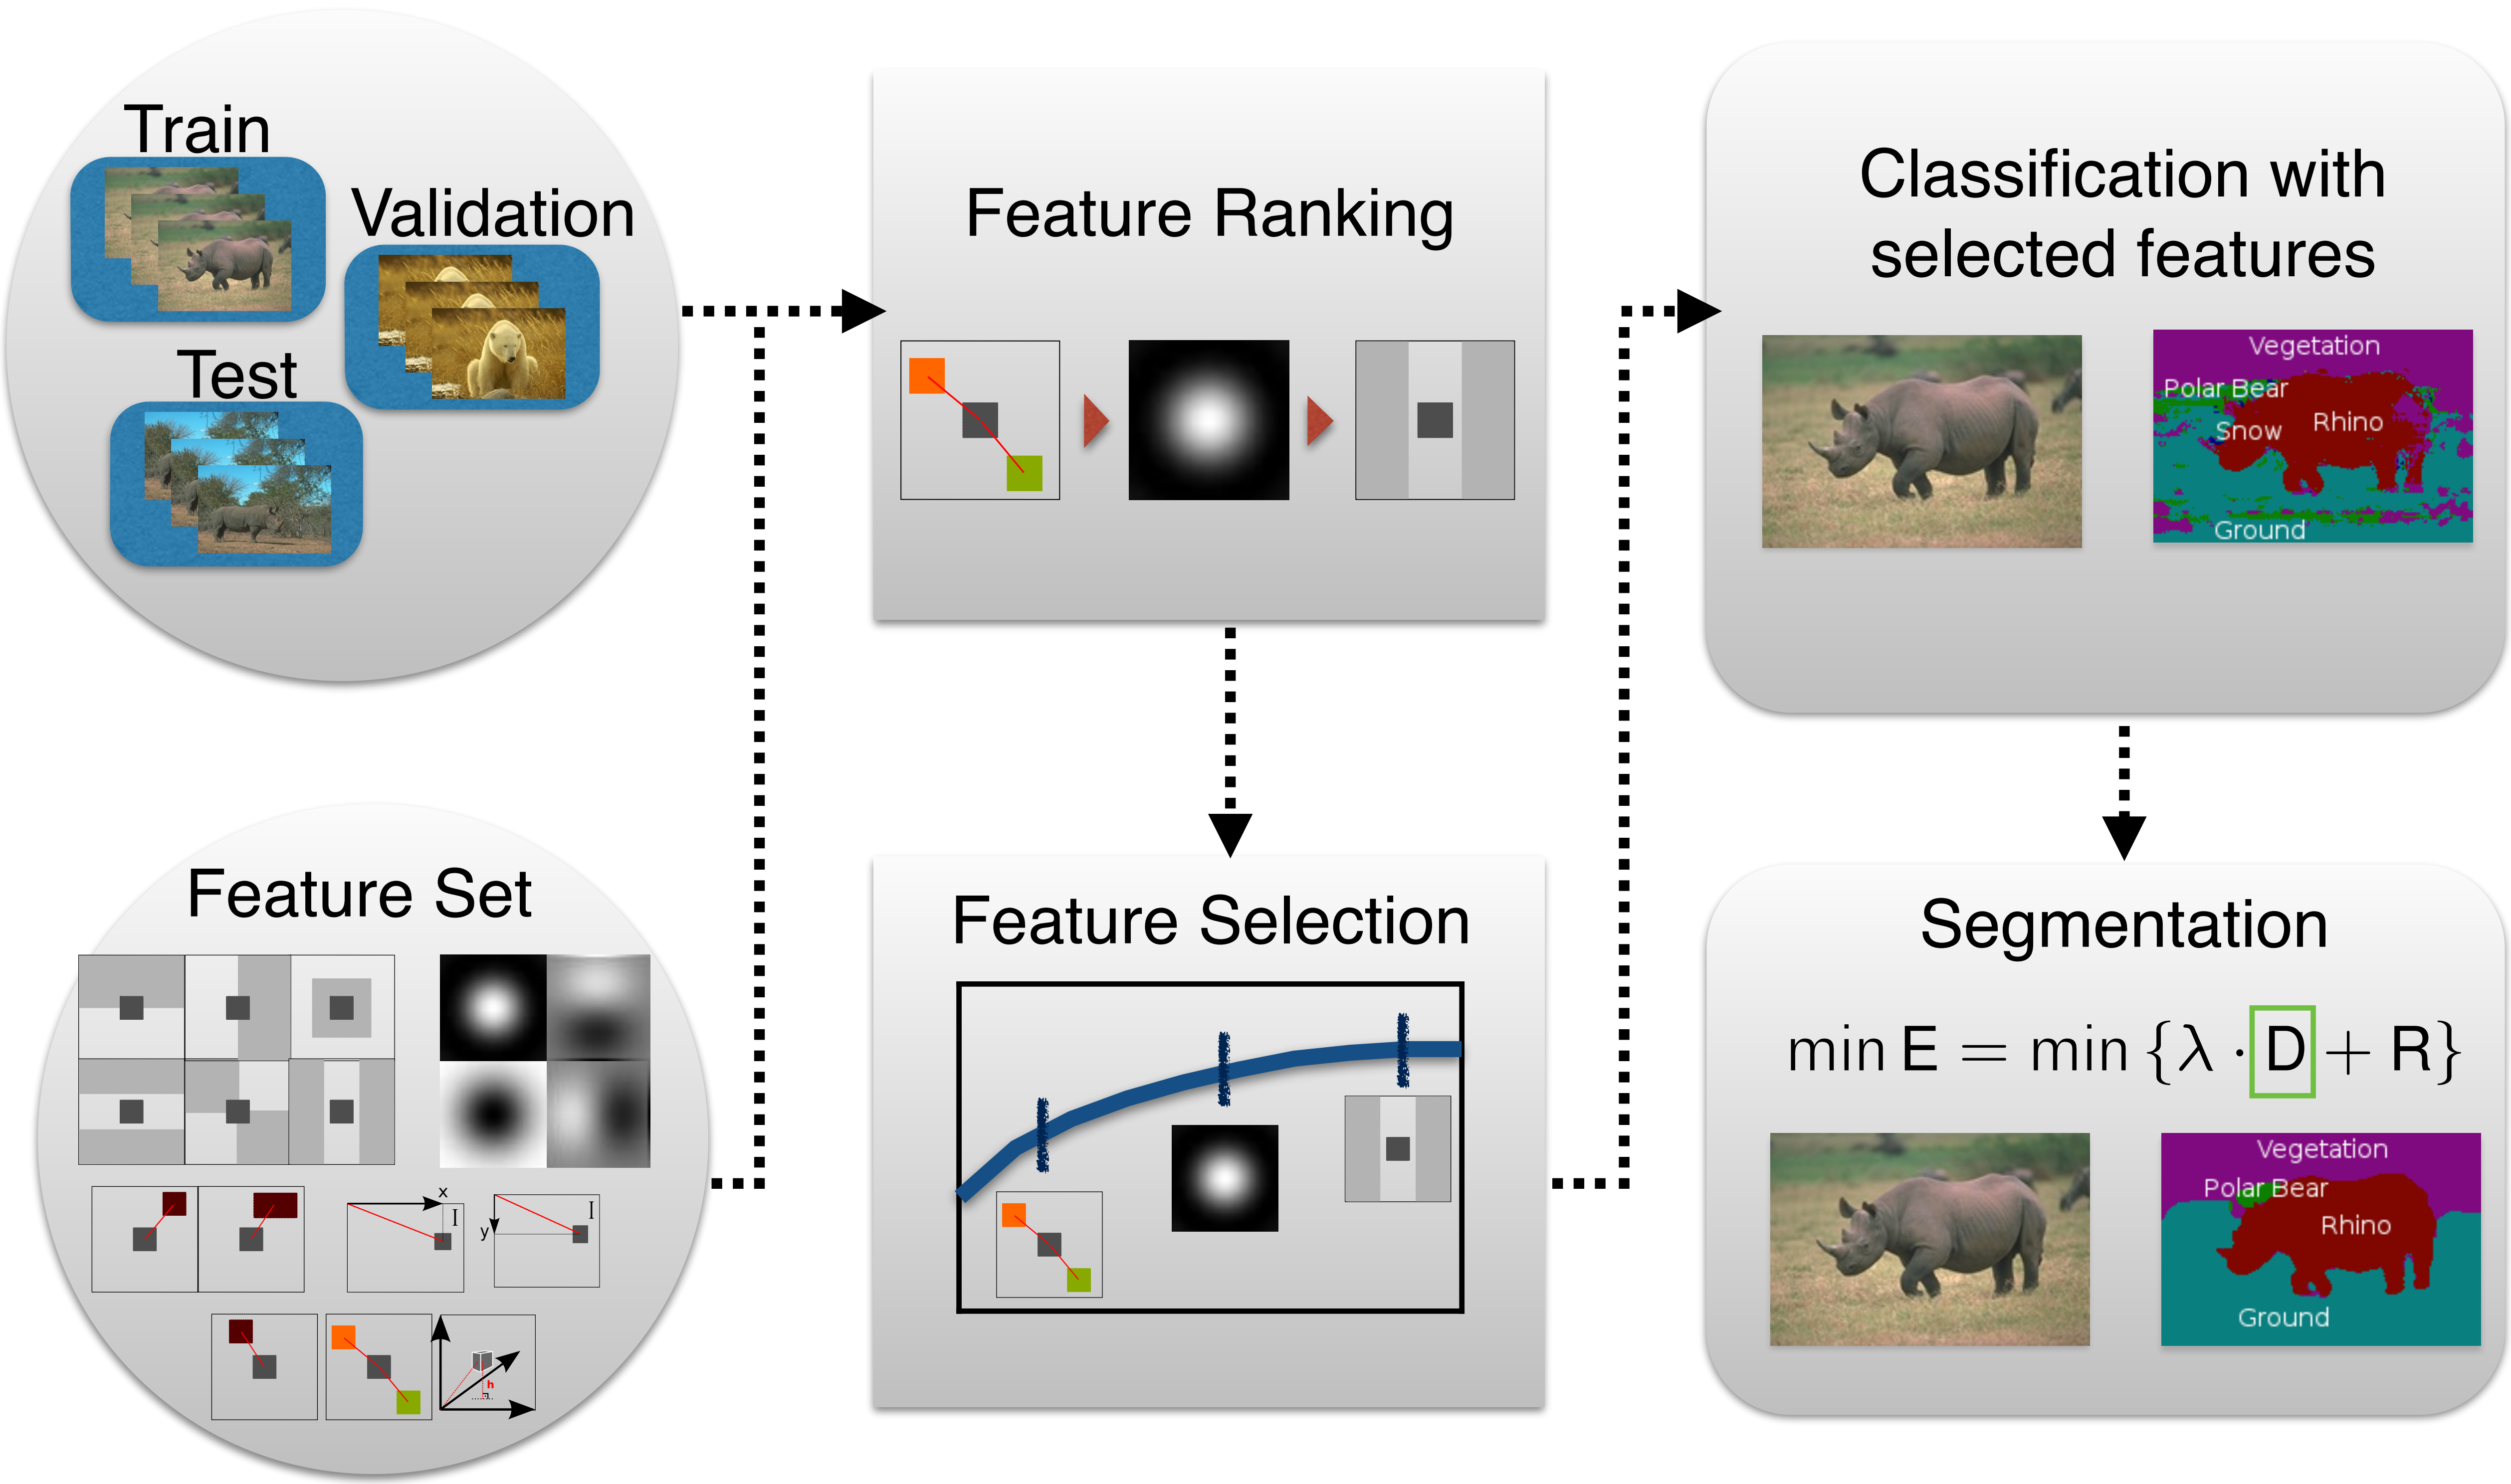

Feature Set

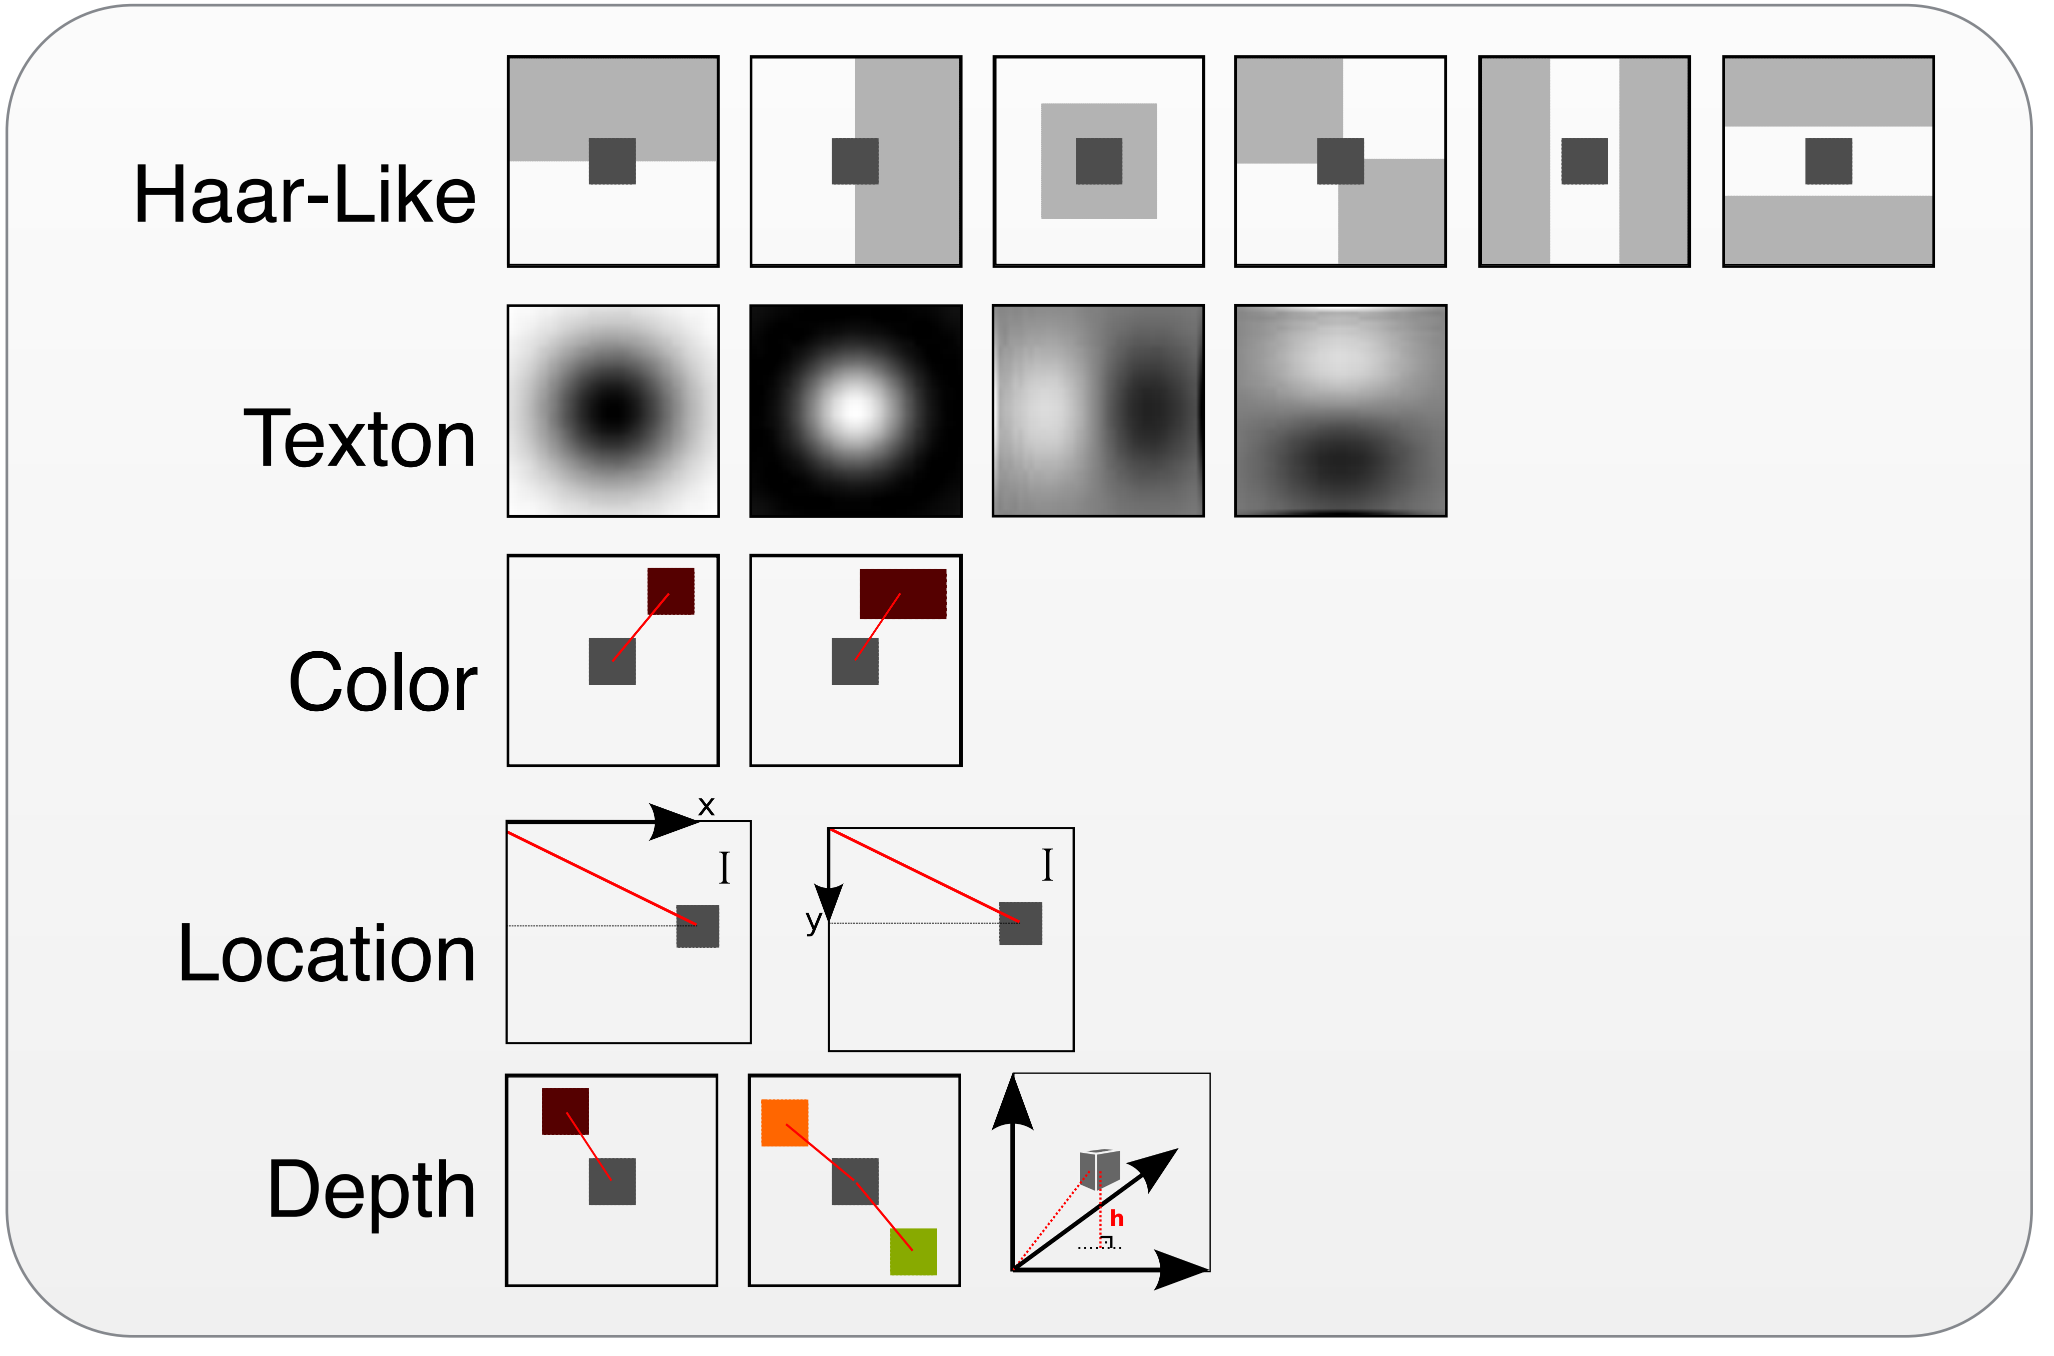

Feature Analysis

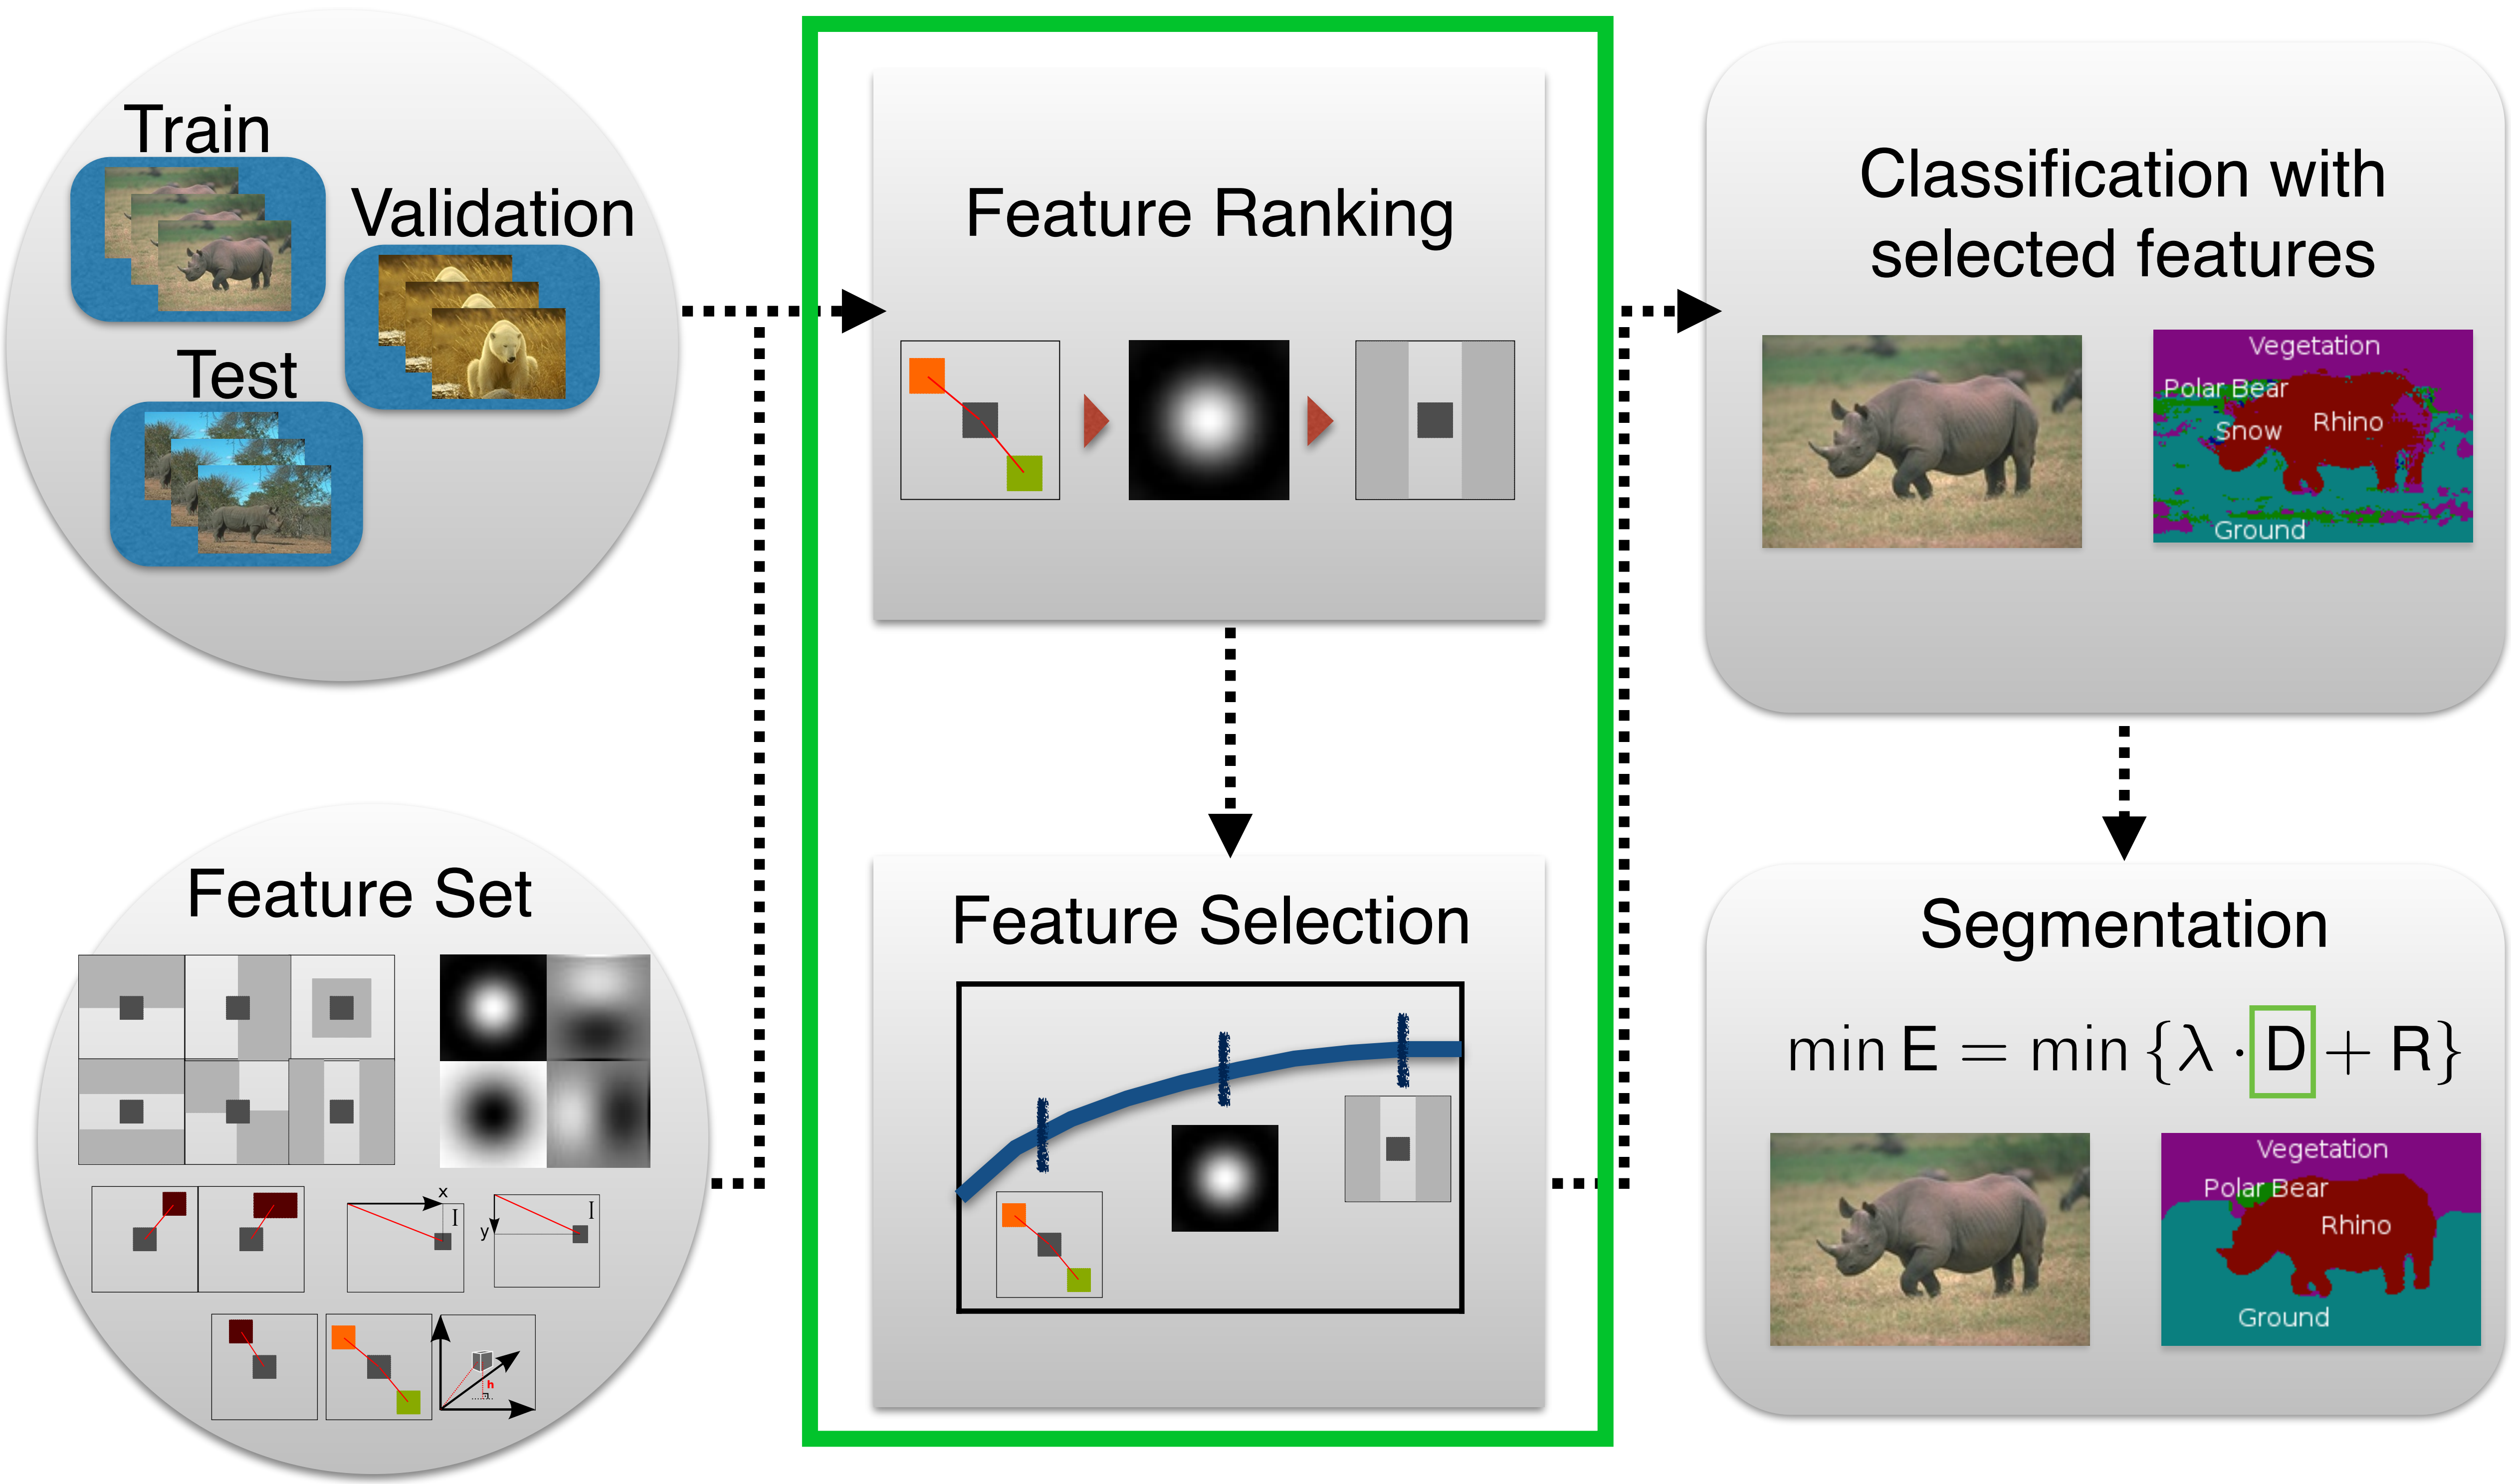

Feature Ranking

Objective:

- *maximize* the relevance between the feature and its class
- *minimize* the redundancy between the feature pairs

$$\max \Phi\left(\text{Rel}, \text{Red}\right), \quad \boxed{\begin{aligned}\Phi &= \text{Relevance} - \text{Redundancy} \\ &= \text{MI}\left(f_i; class\right) - \frac{1}{m-1}\sum_{i \neq j} \text{MI}\left(f_i; f_j\right)\end{aligned}}$$

$$\text{MI}\left(X;Y\right) = \int_Y \int_X \log \frac{p\left(x,y\right)}{p\left(x\right)p\left(y\right)} dxdy$$

Ranking:

- iteratively rank the features, select one feature at a time
- maximize the objective function at each iteration:

$$f_m = \underset{f_i \in \mathcal{F} \setminus \mathcal{F}_{m-1}}{\arg\max} \left[\text{MI}\left(f_i; c\right) - \frac{1}{m-1} \sum_{f_j \in \mathcal{F}_{m-1}} \text{MI}\left(f_i; f_j\right) \right]$$

Incremental Feature Analysis

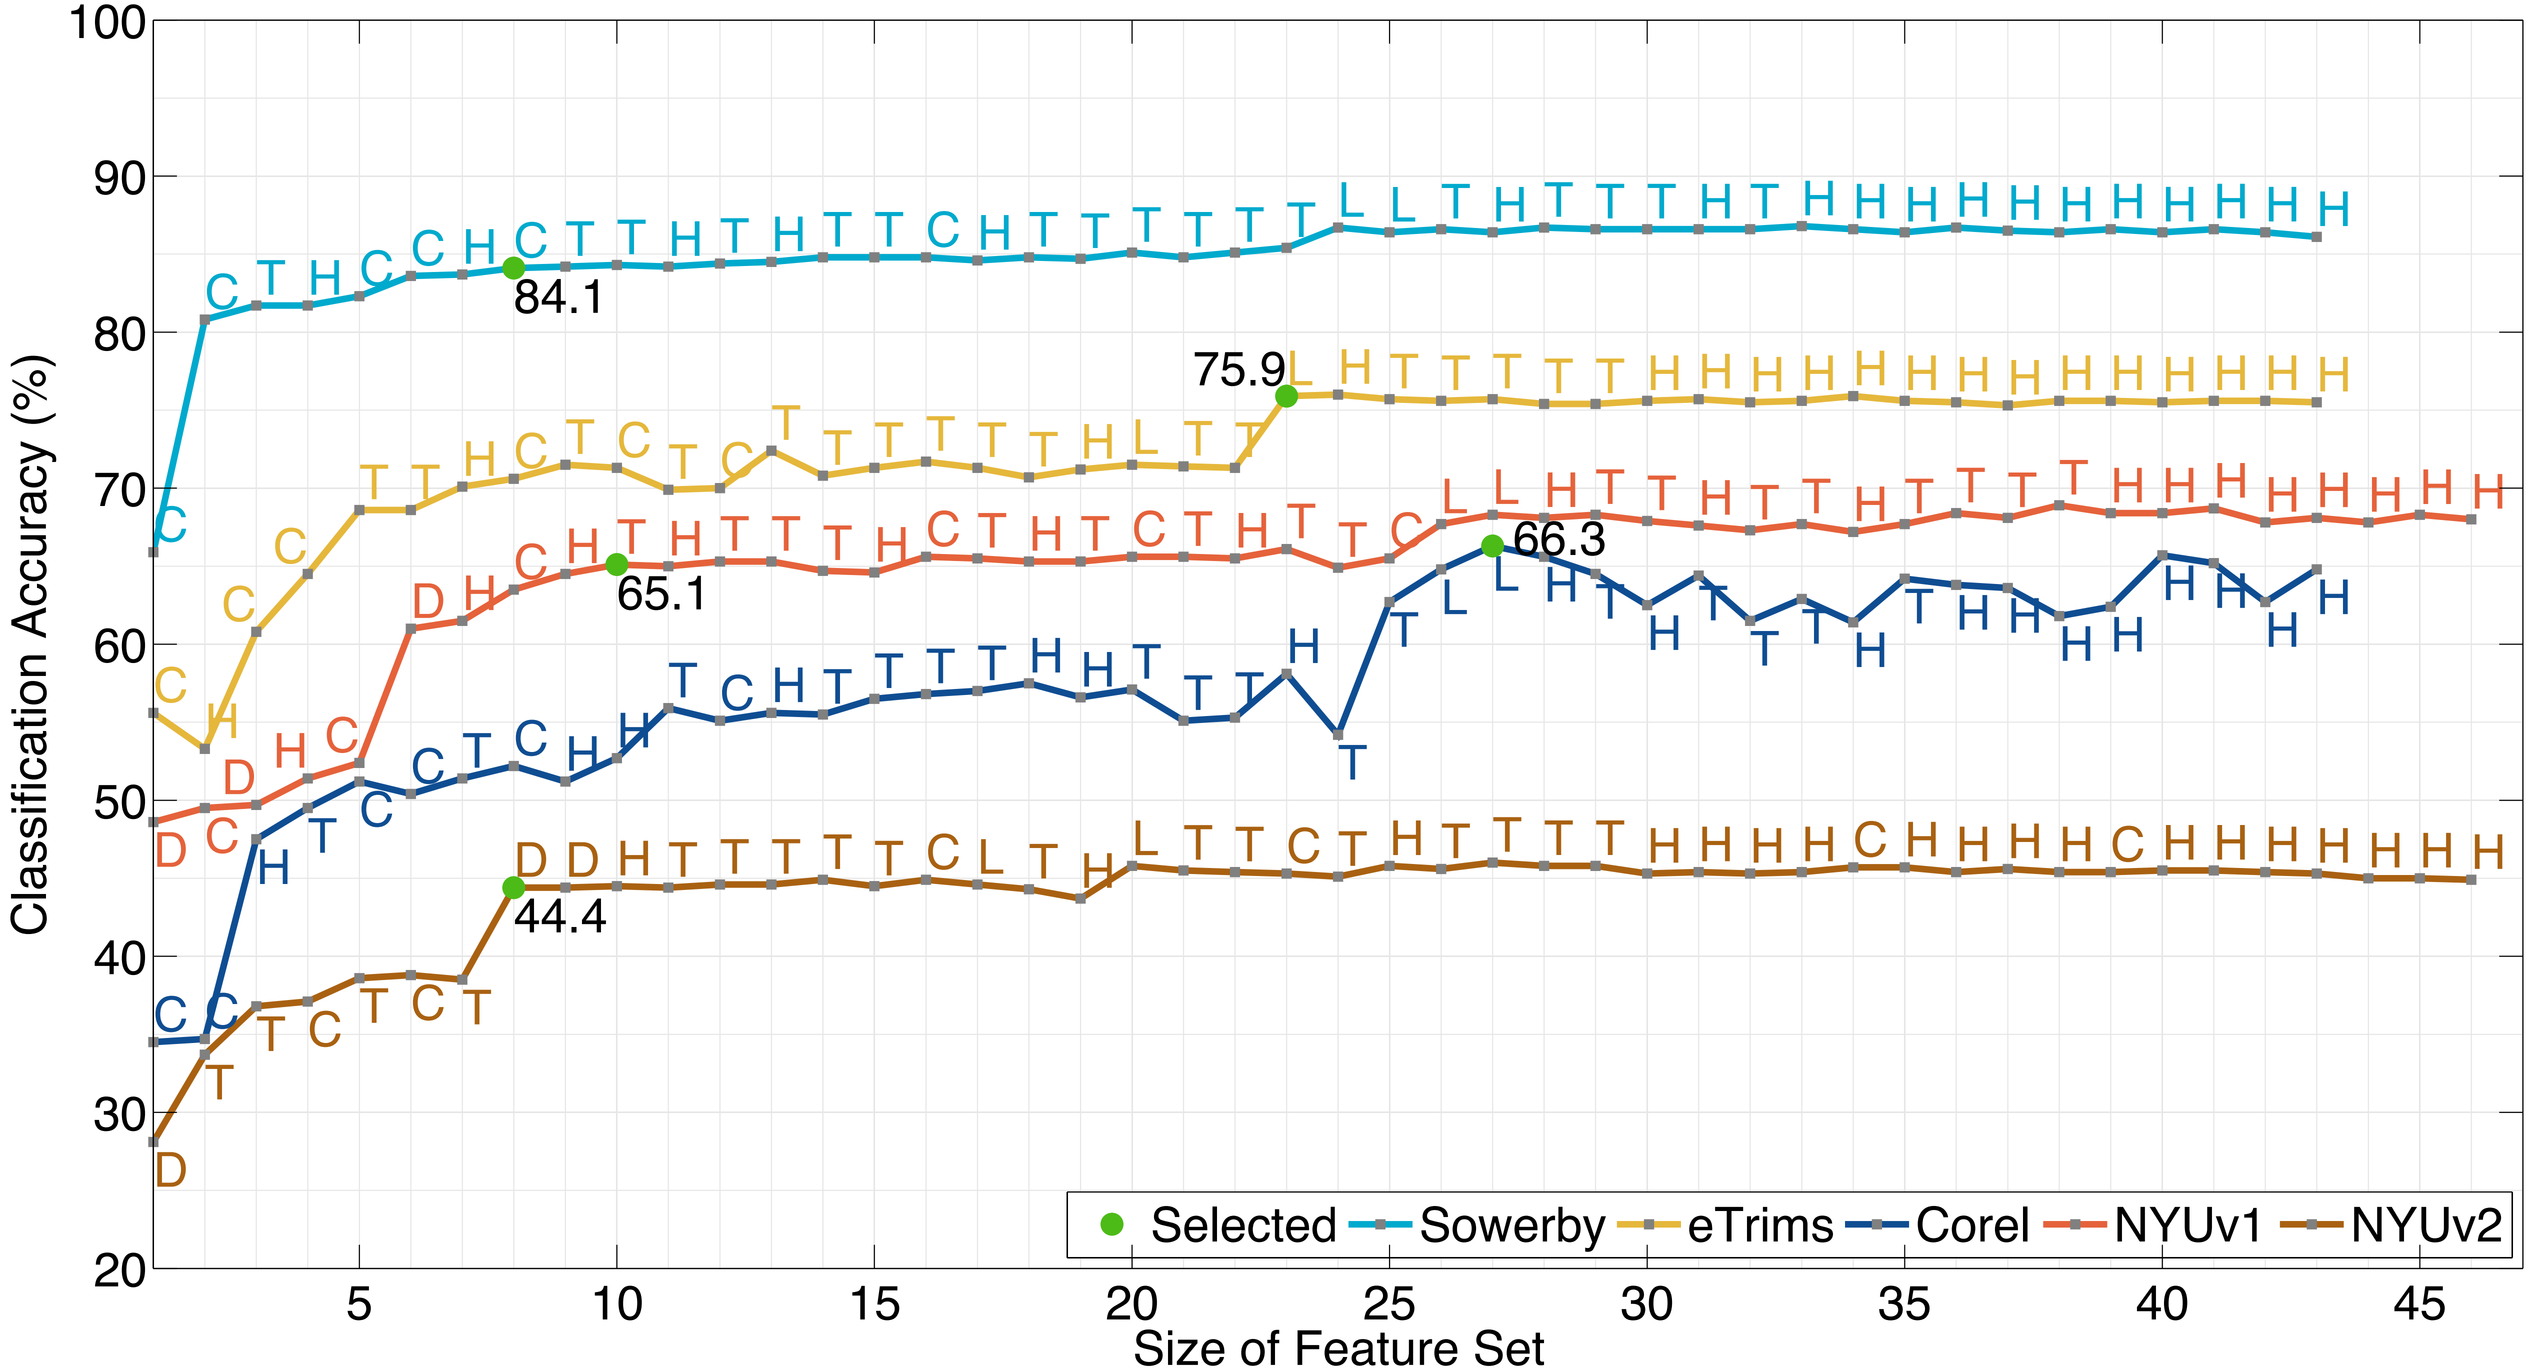

Relevance of **H**aar-Like Features

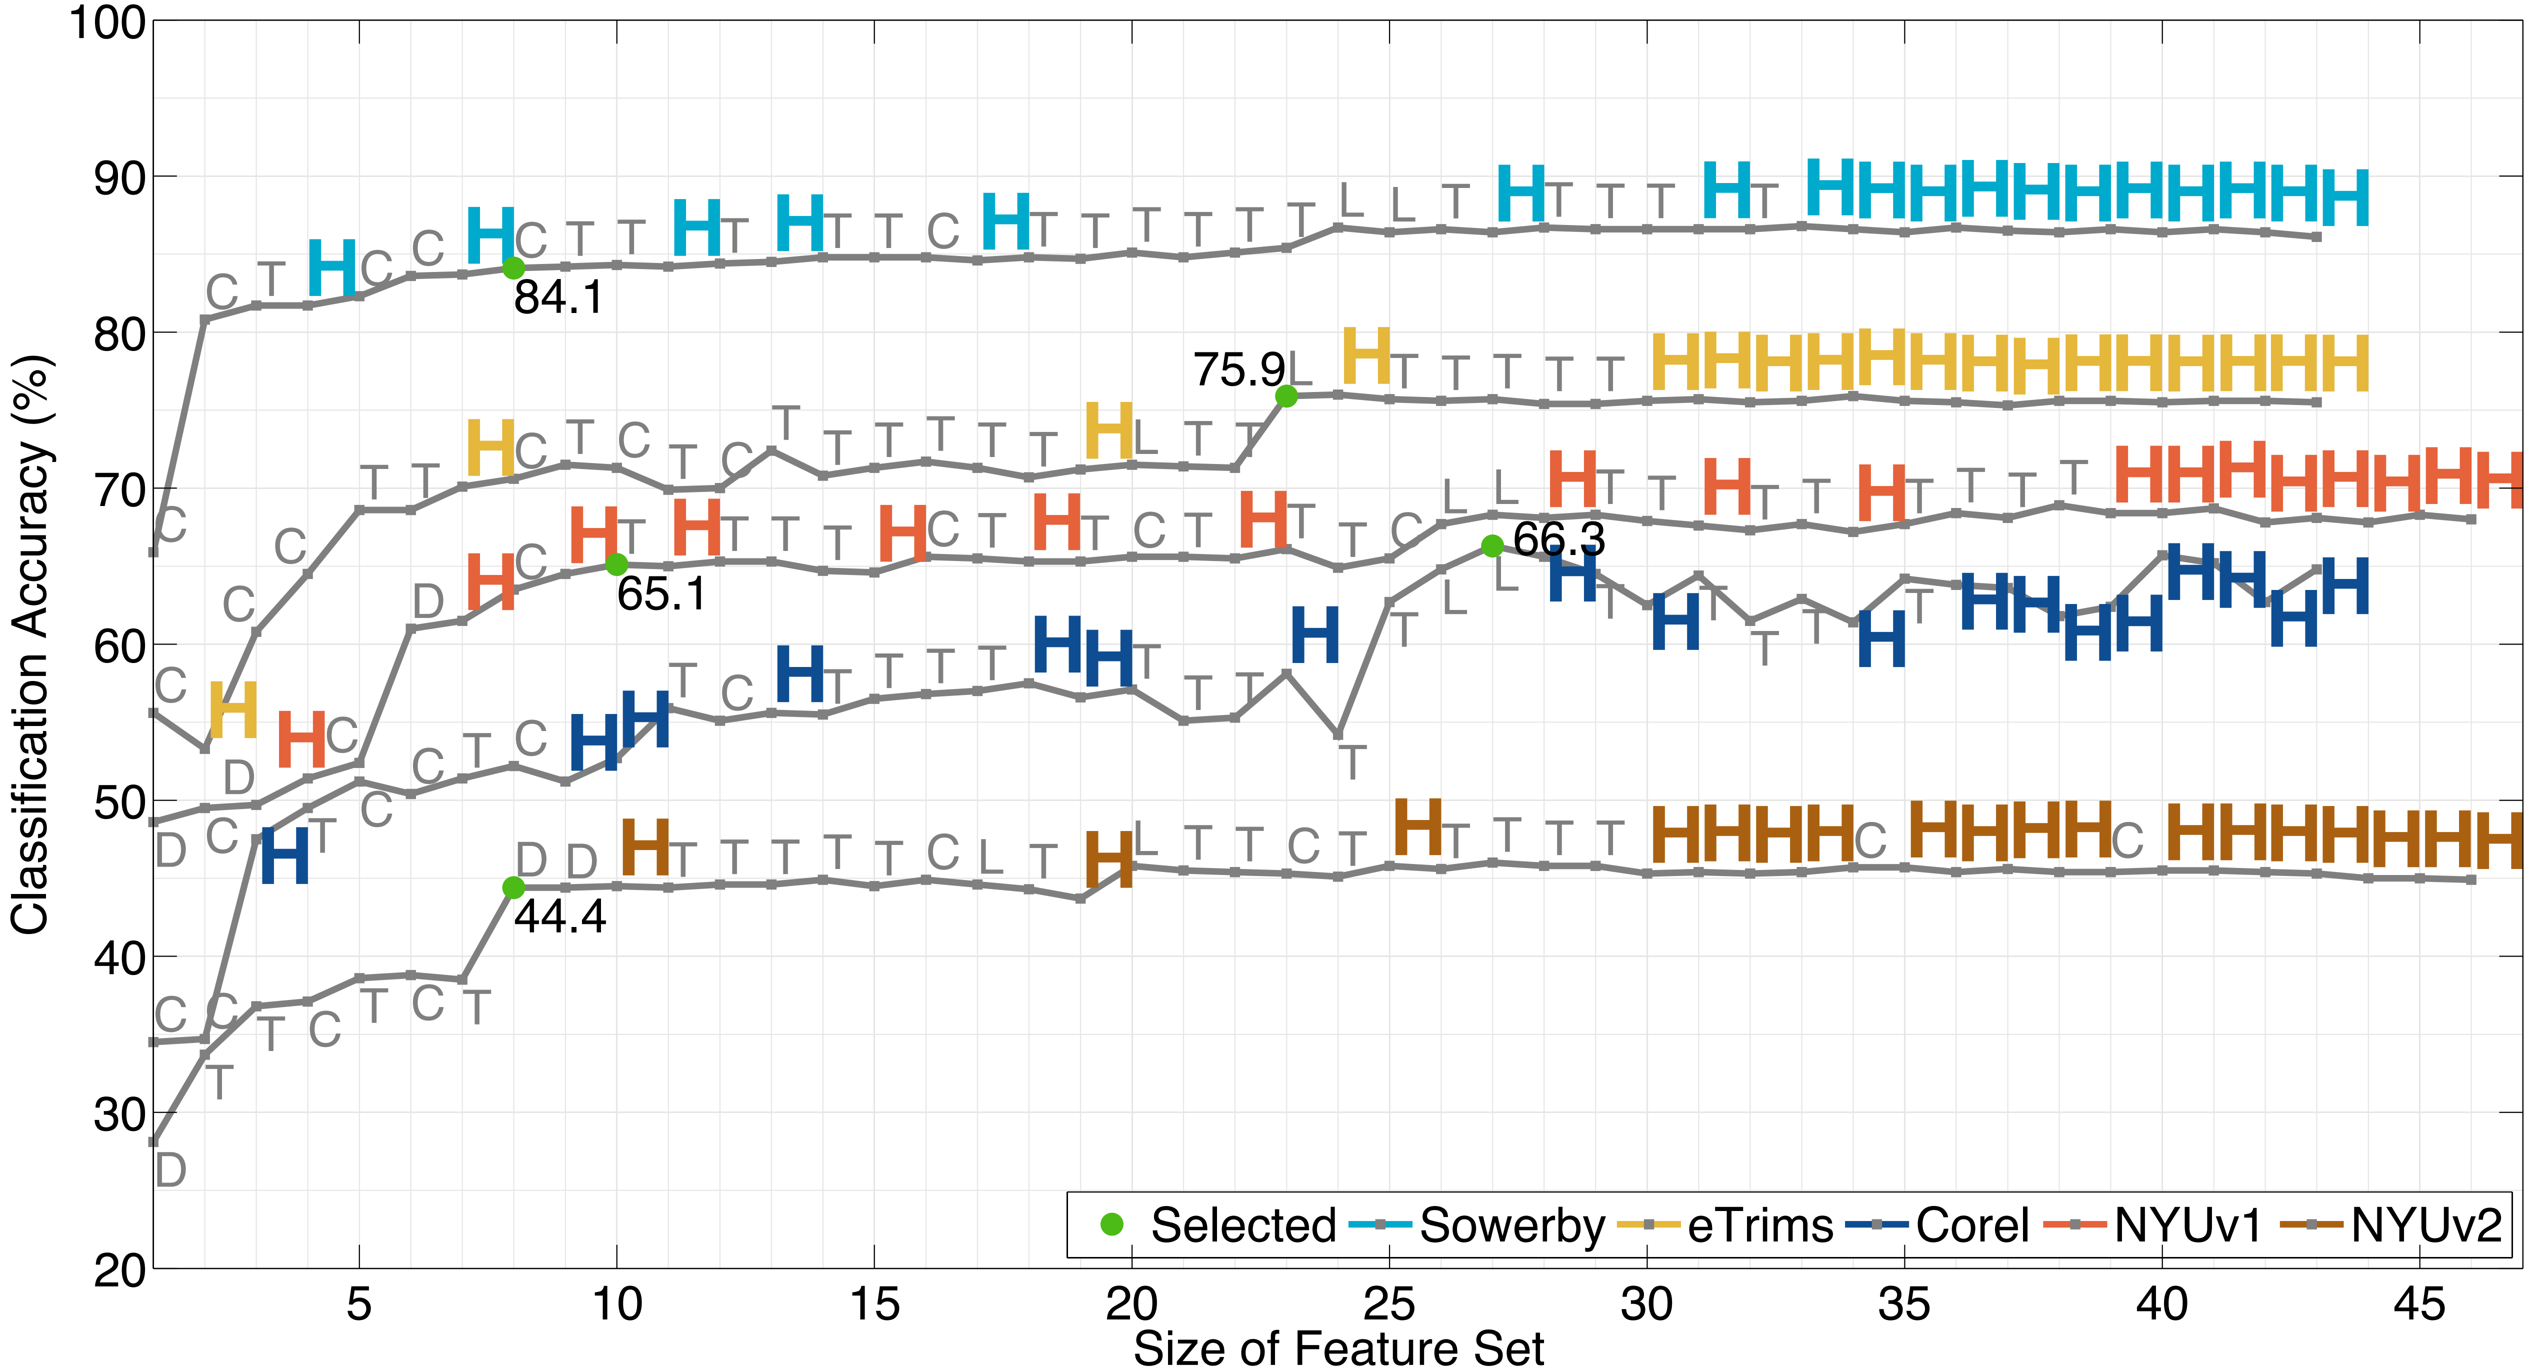

Relevance of **T**exton Features

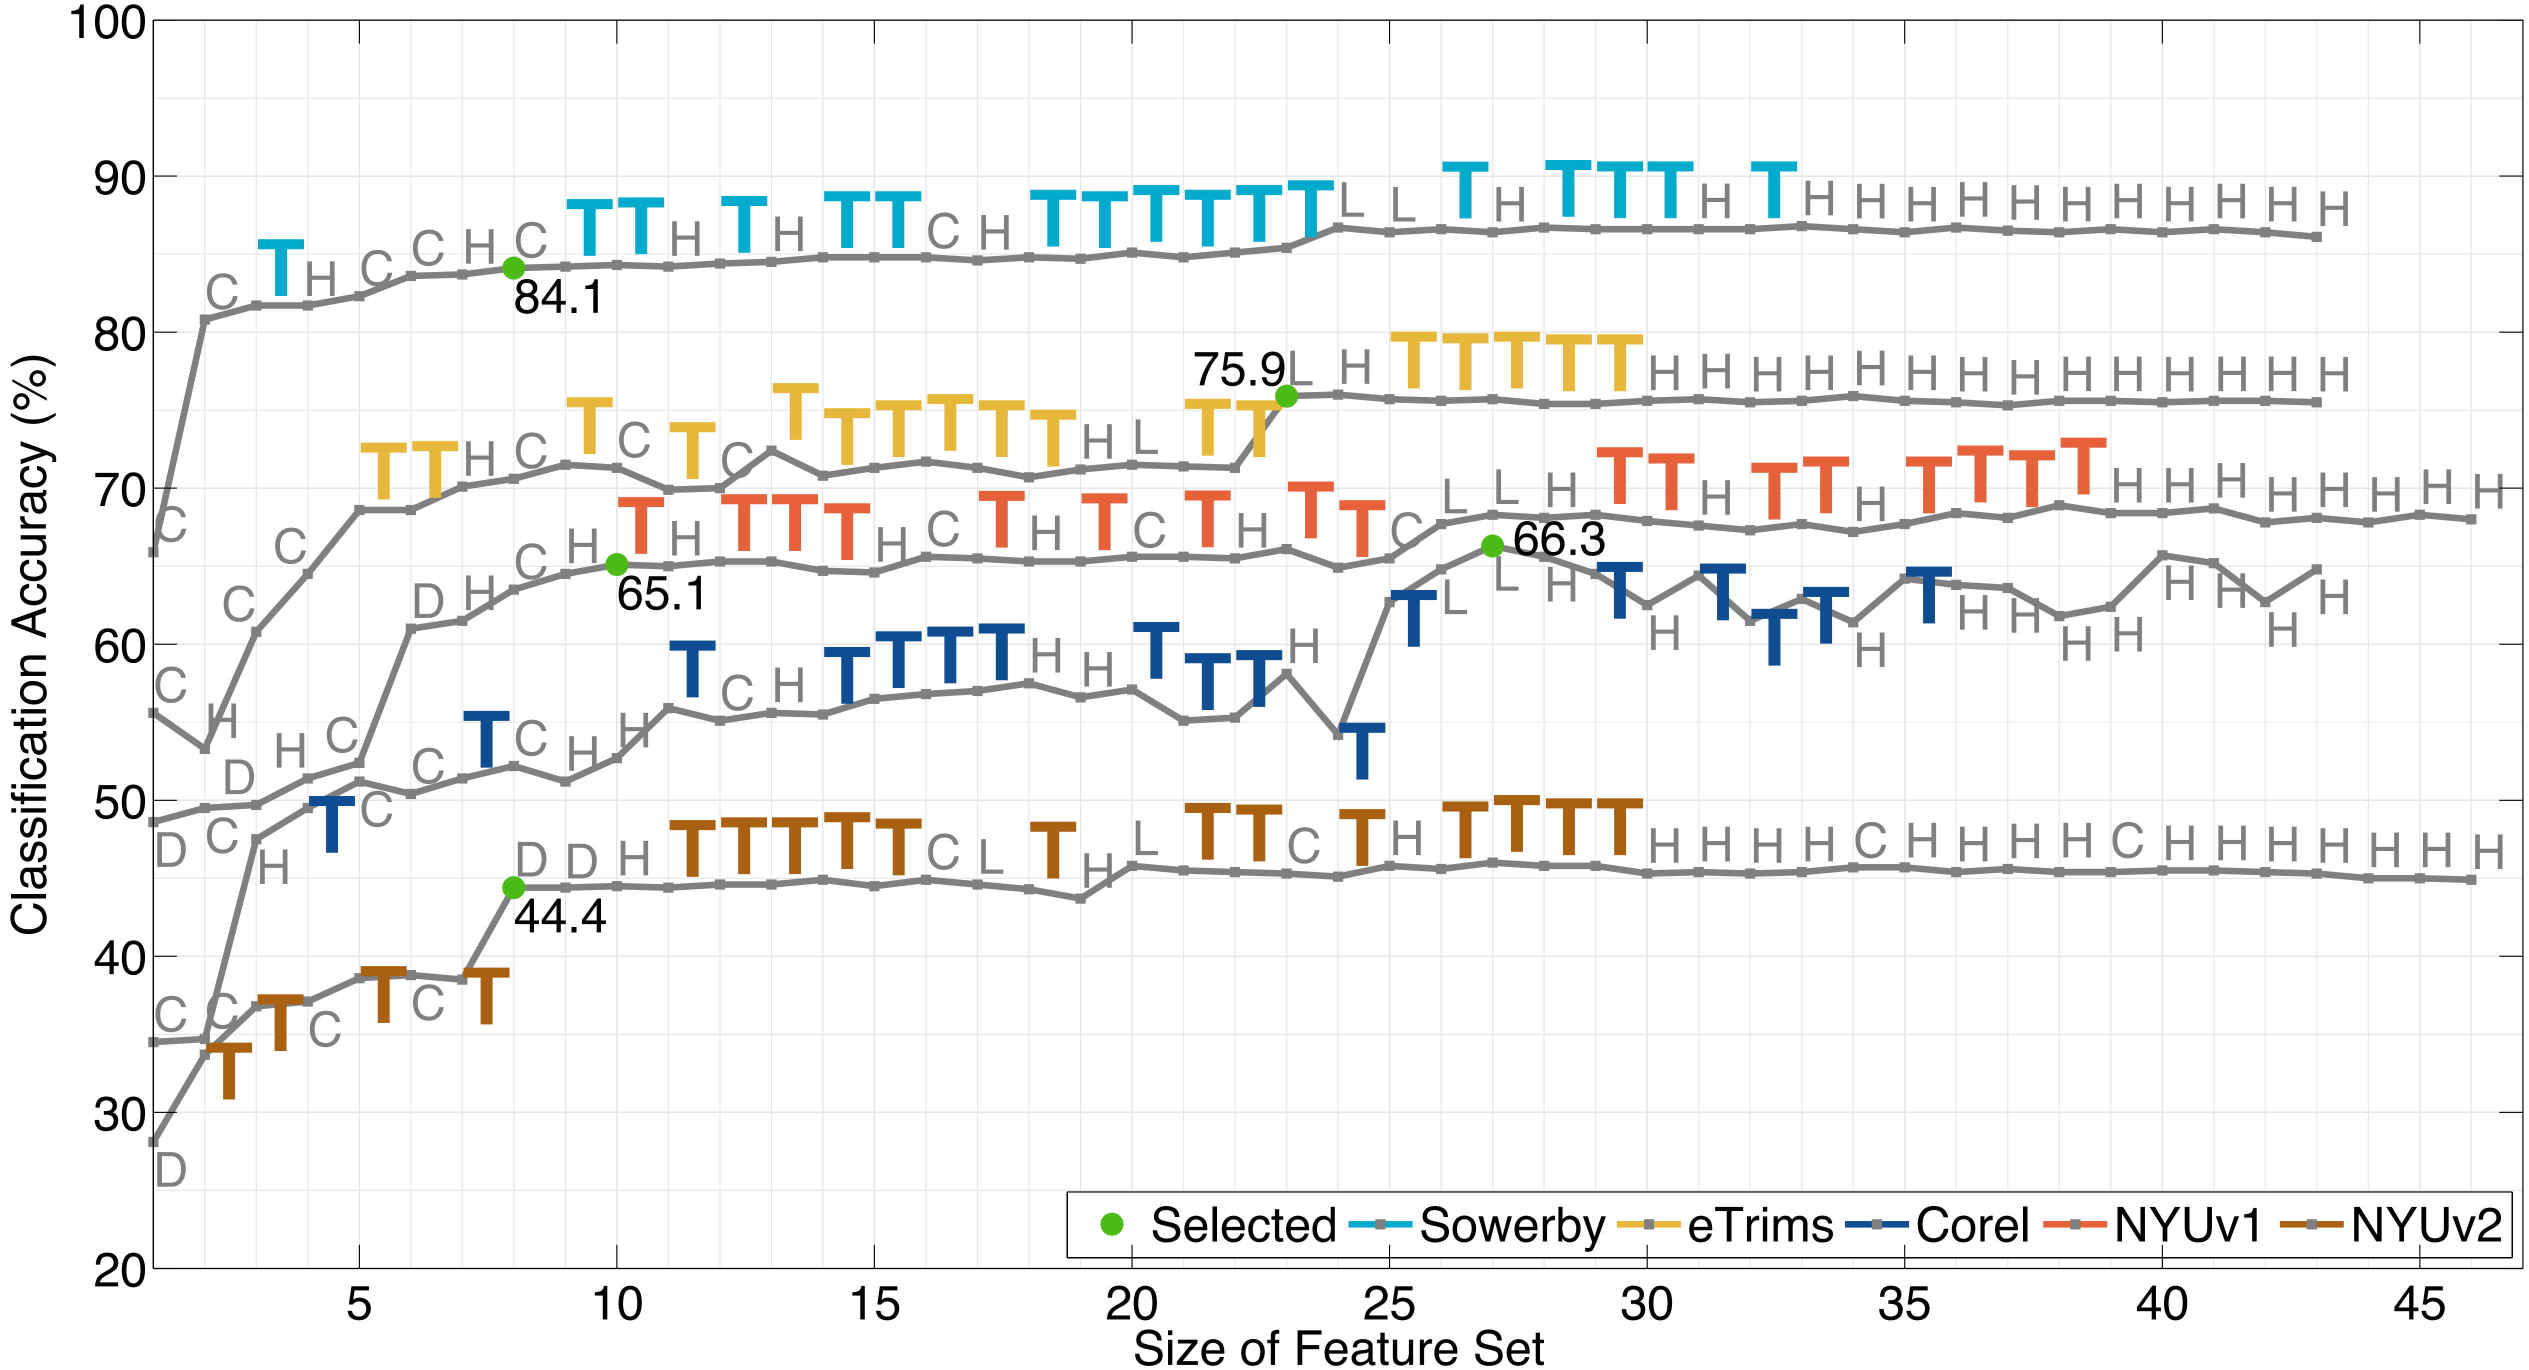

Relevance of **C**olor Features

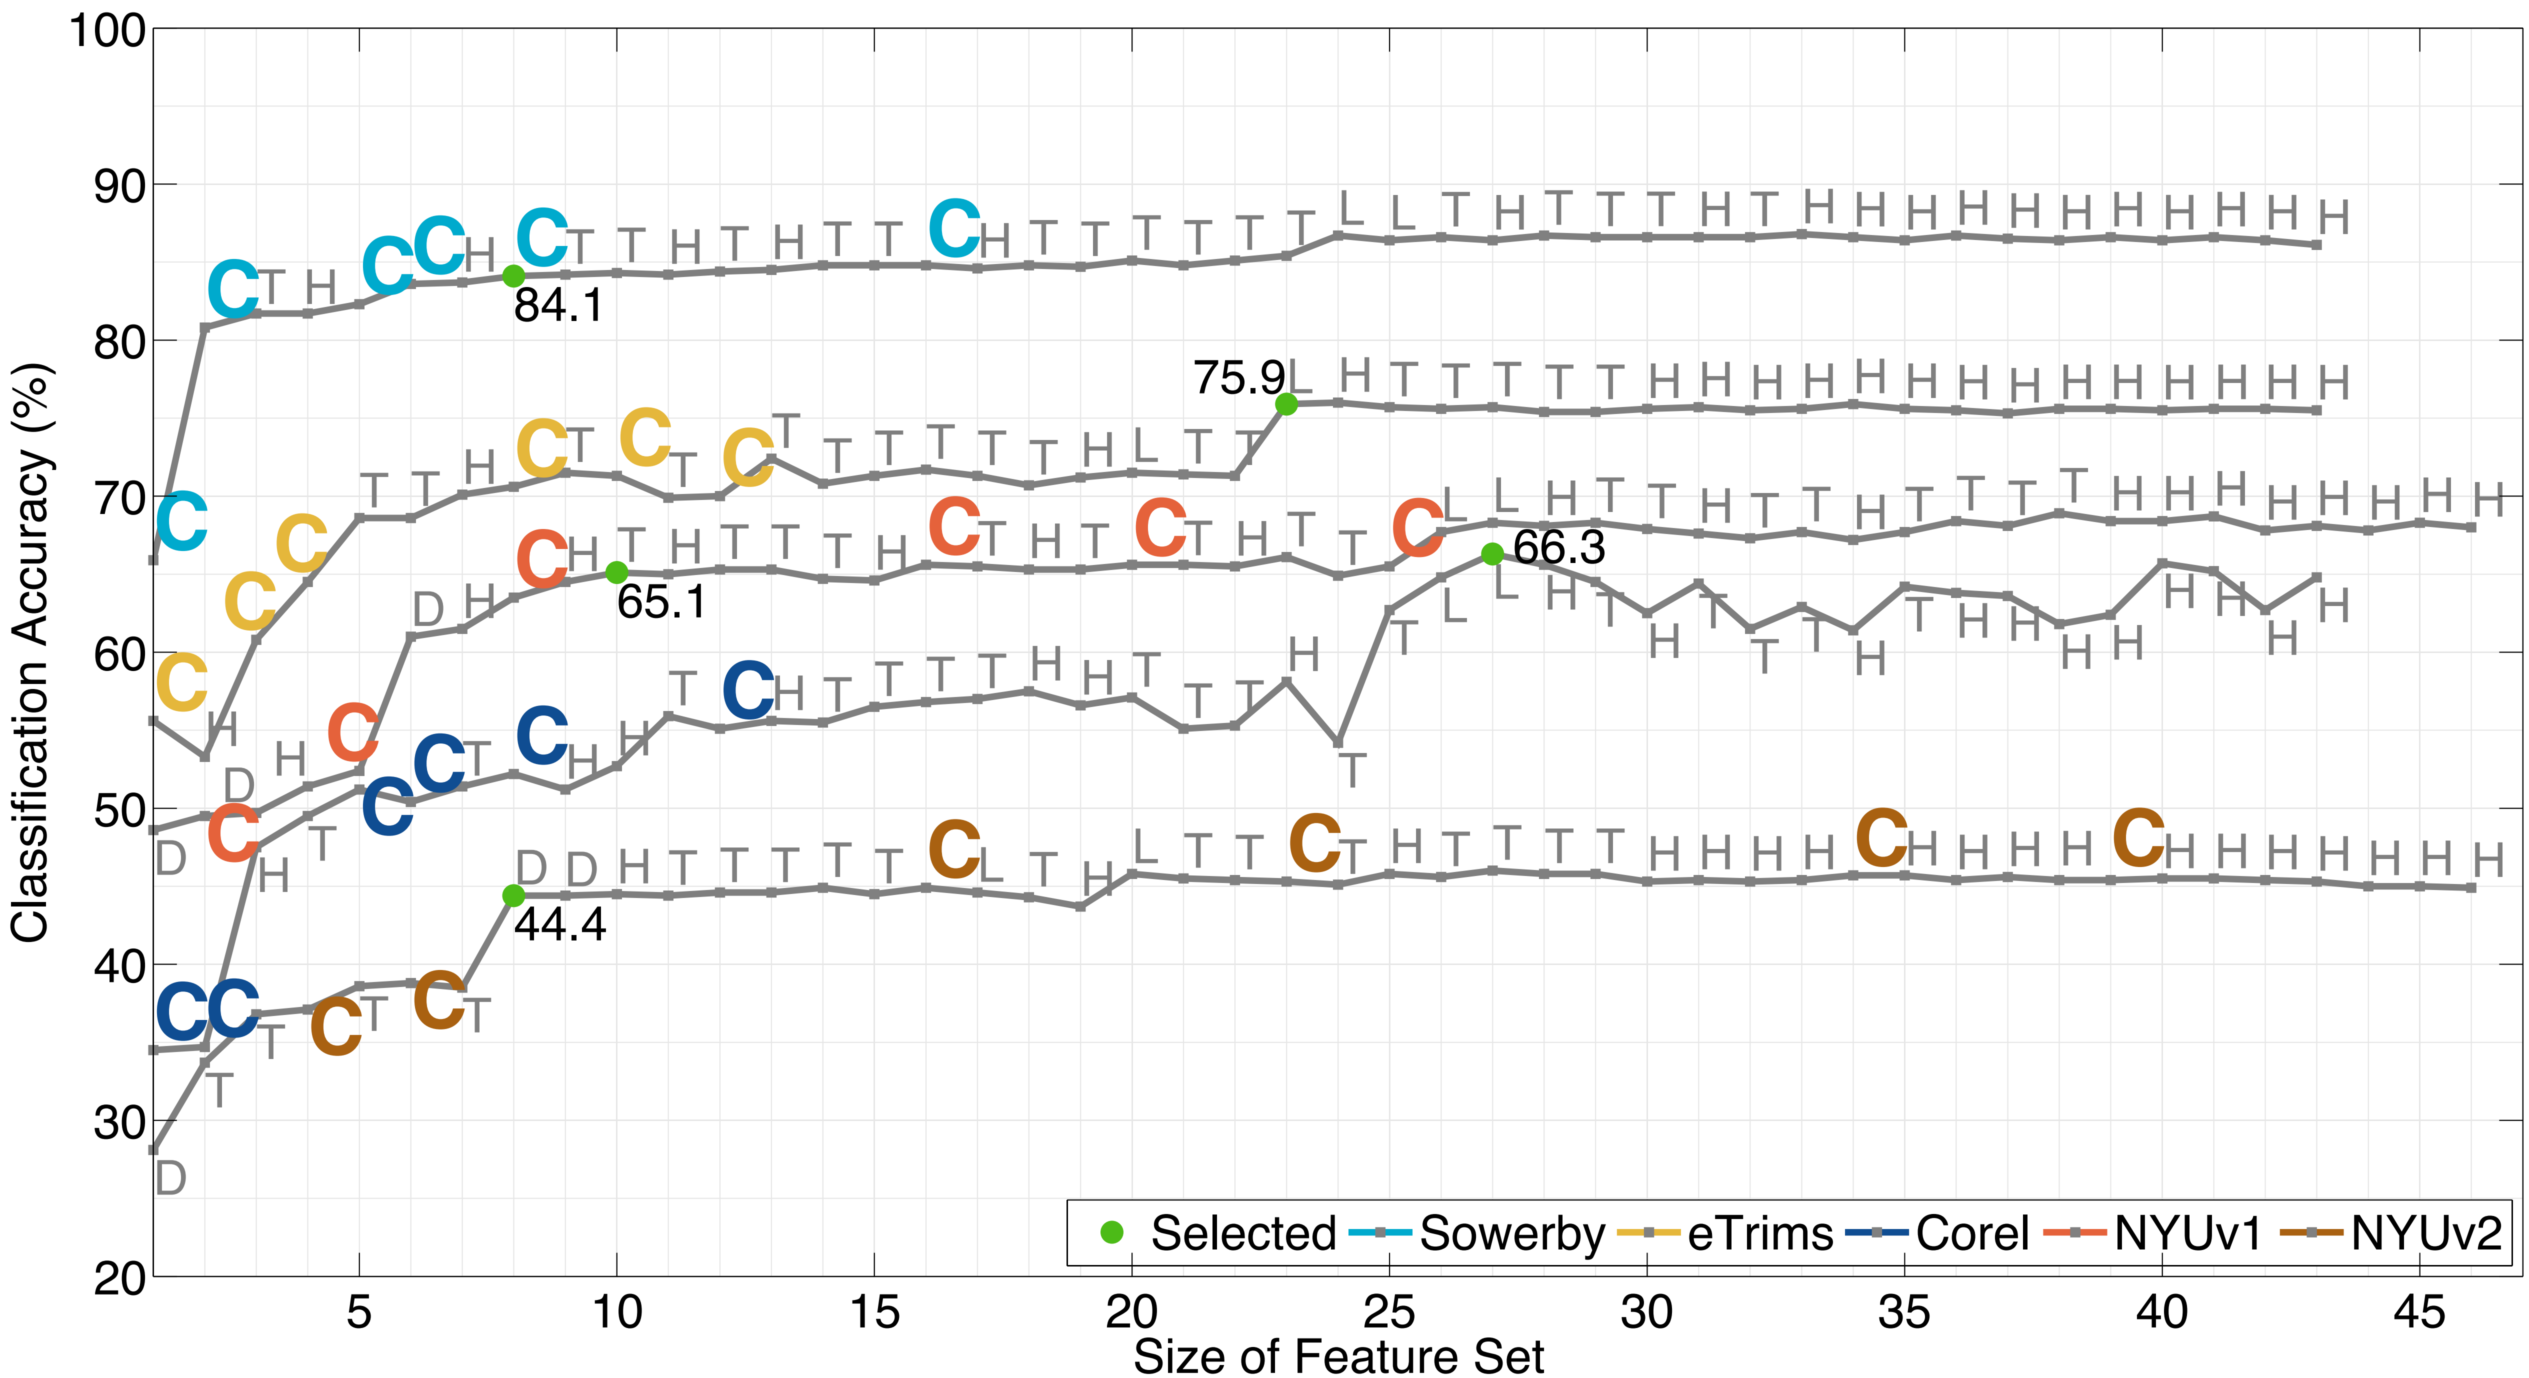

Relevance of **L**ocation Features

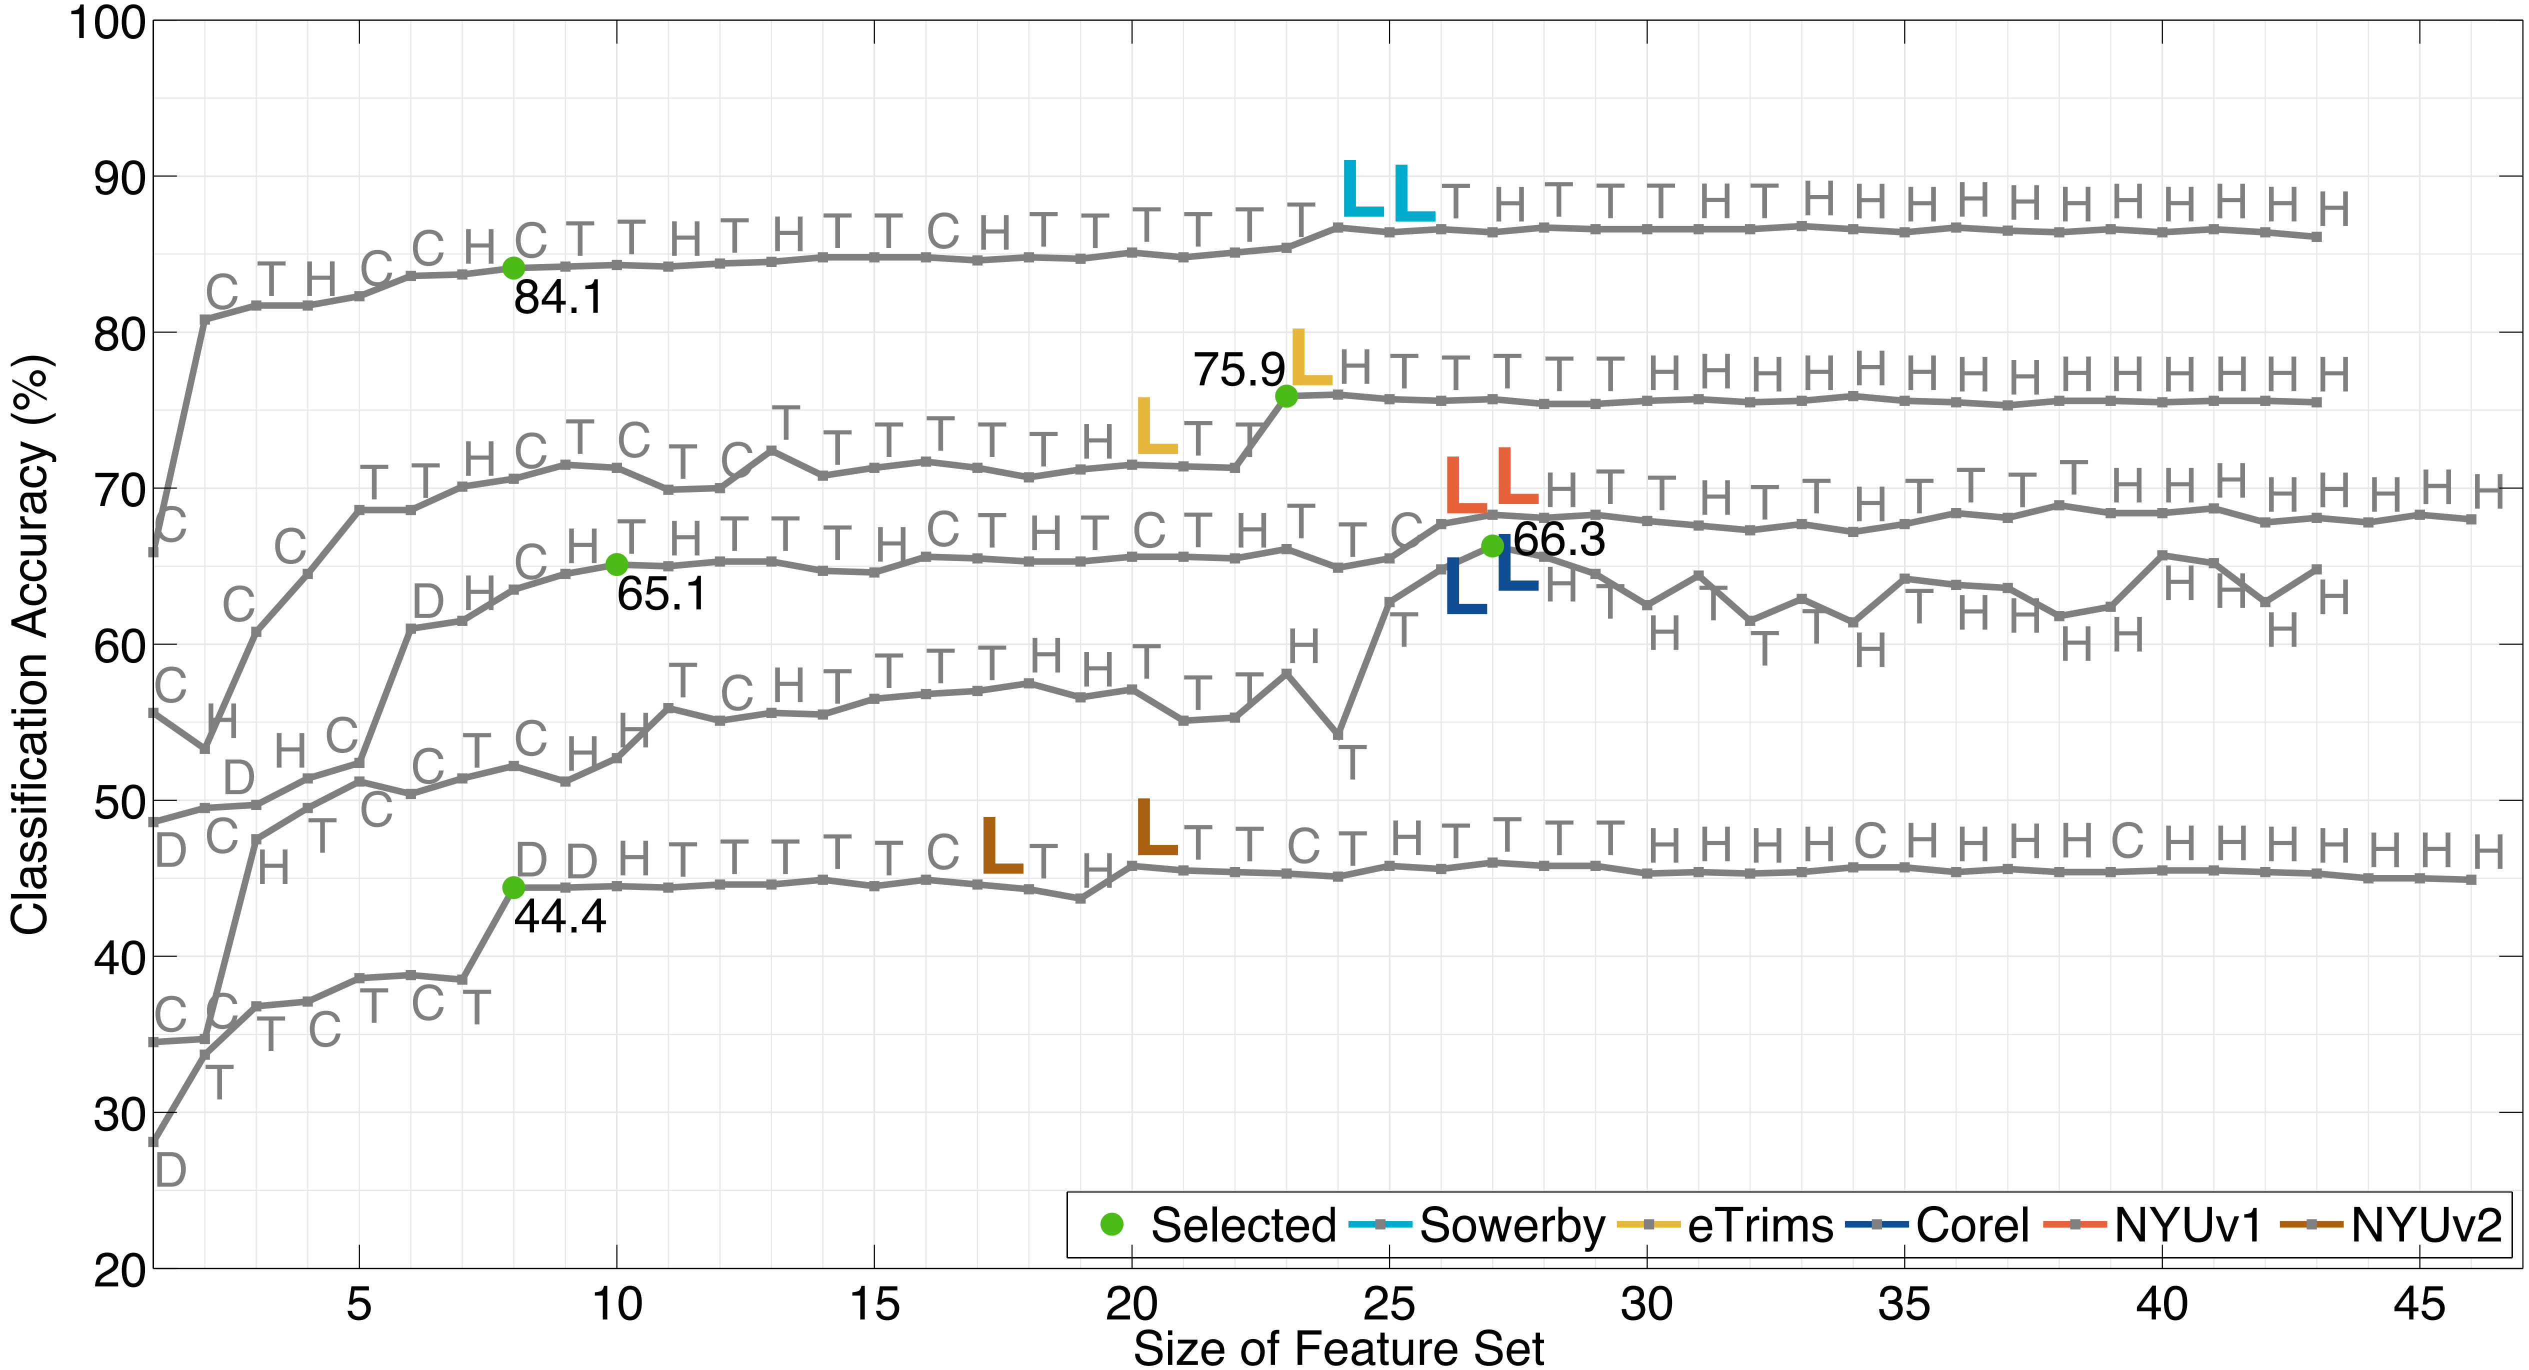

Relevance of **D**epth Features

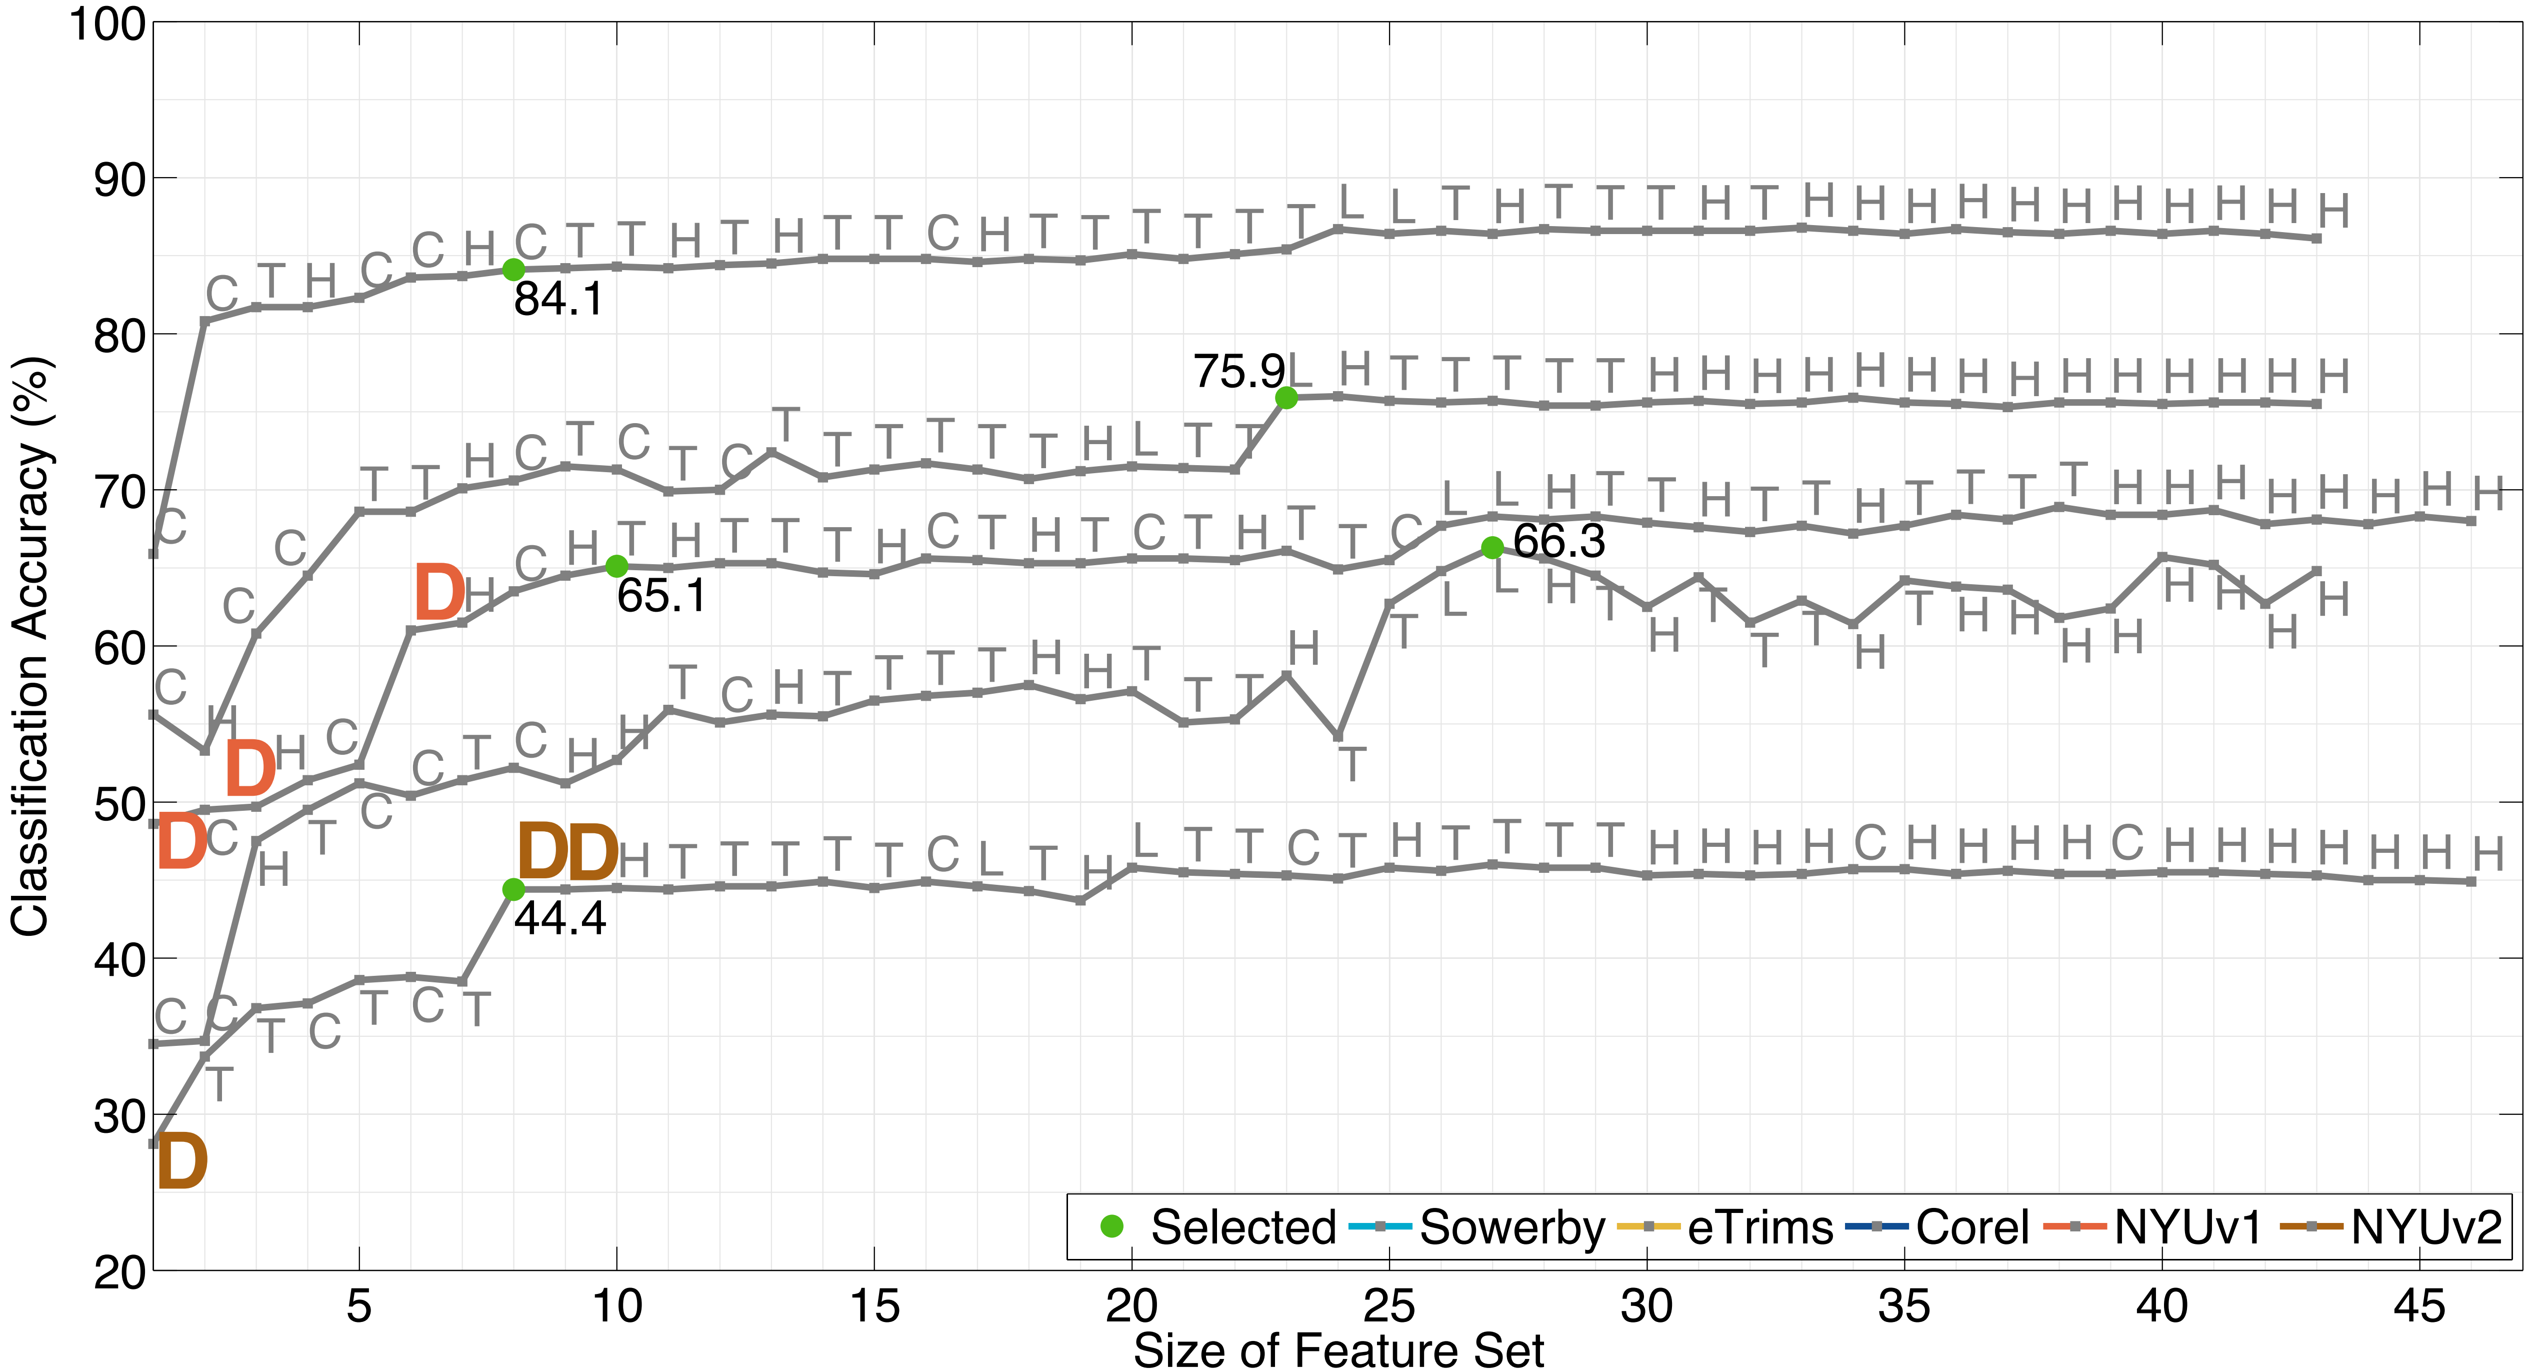

Feature Selection

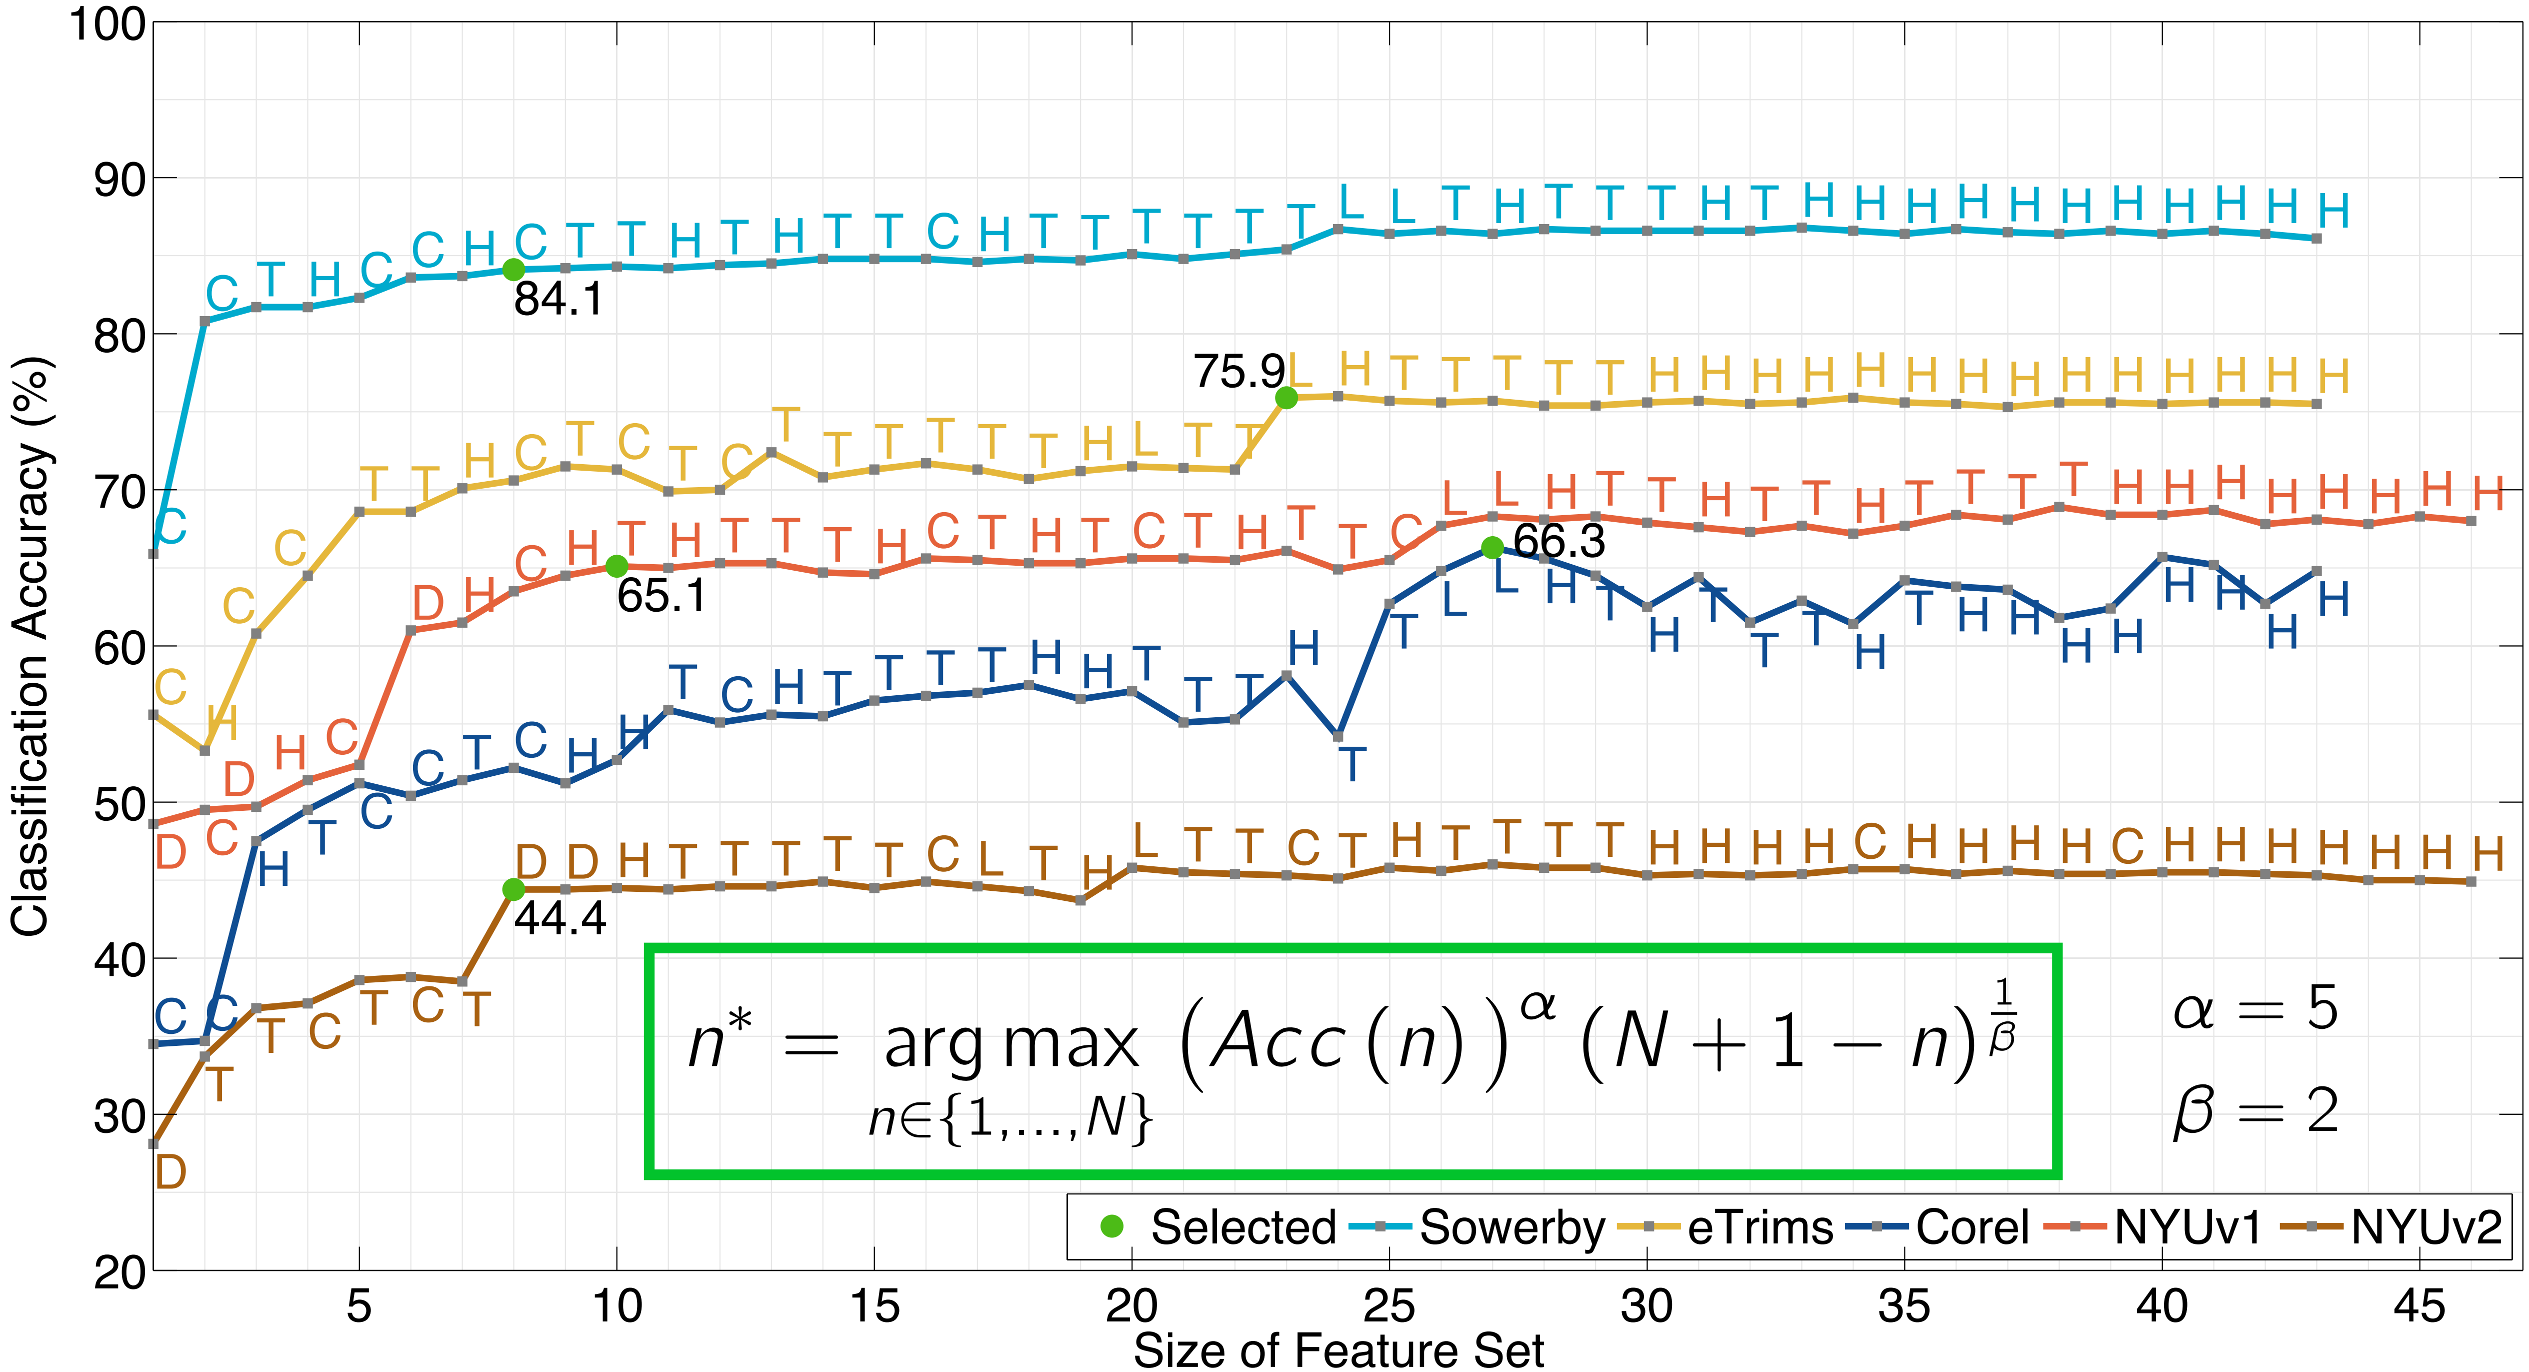

Number of Selected Features

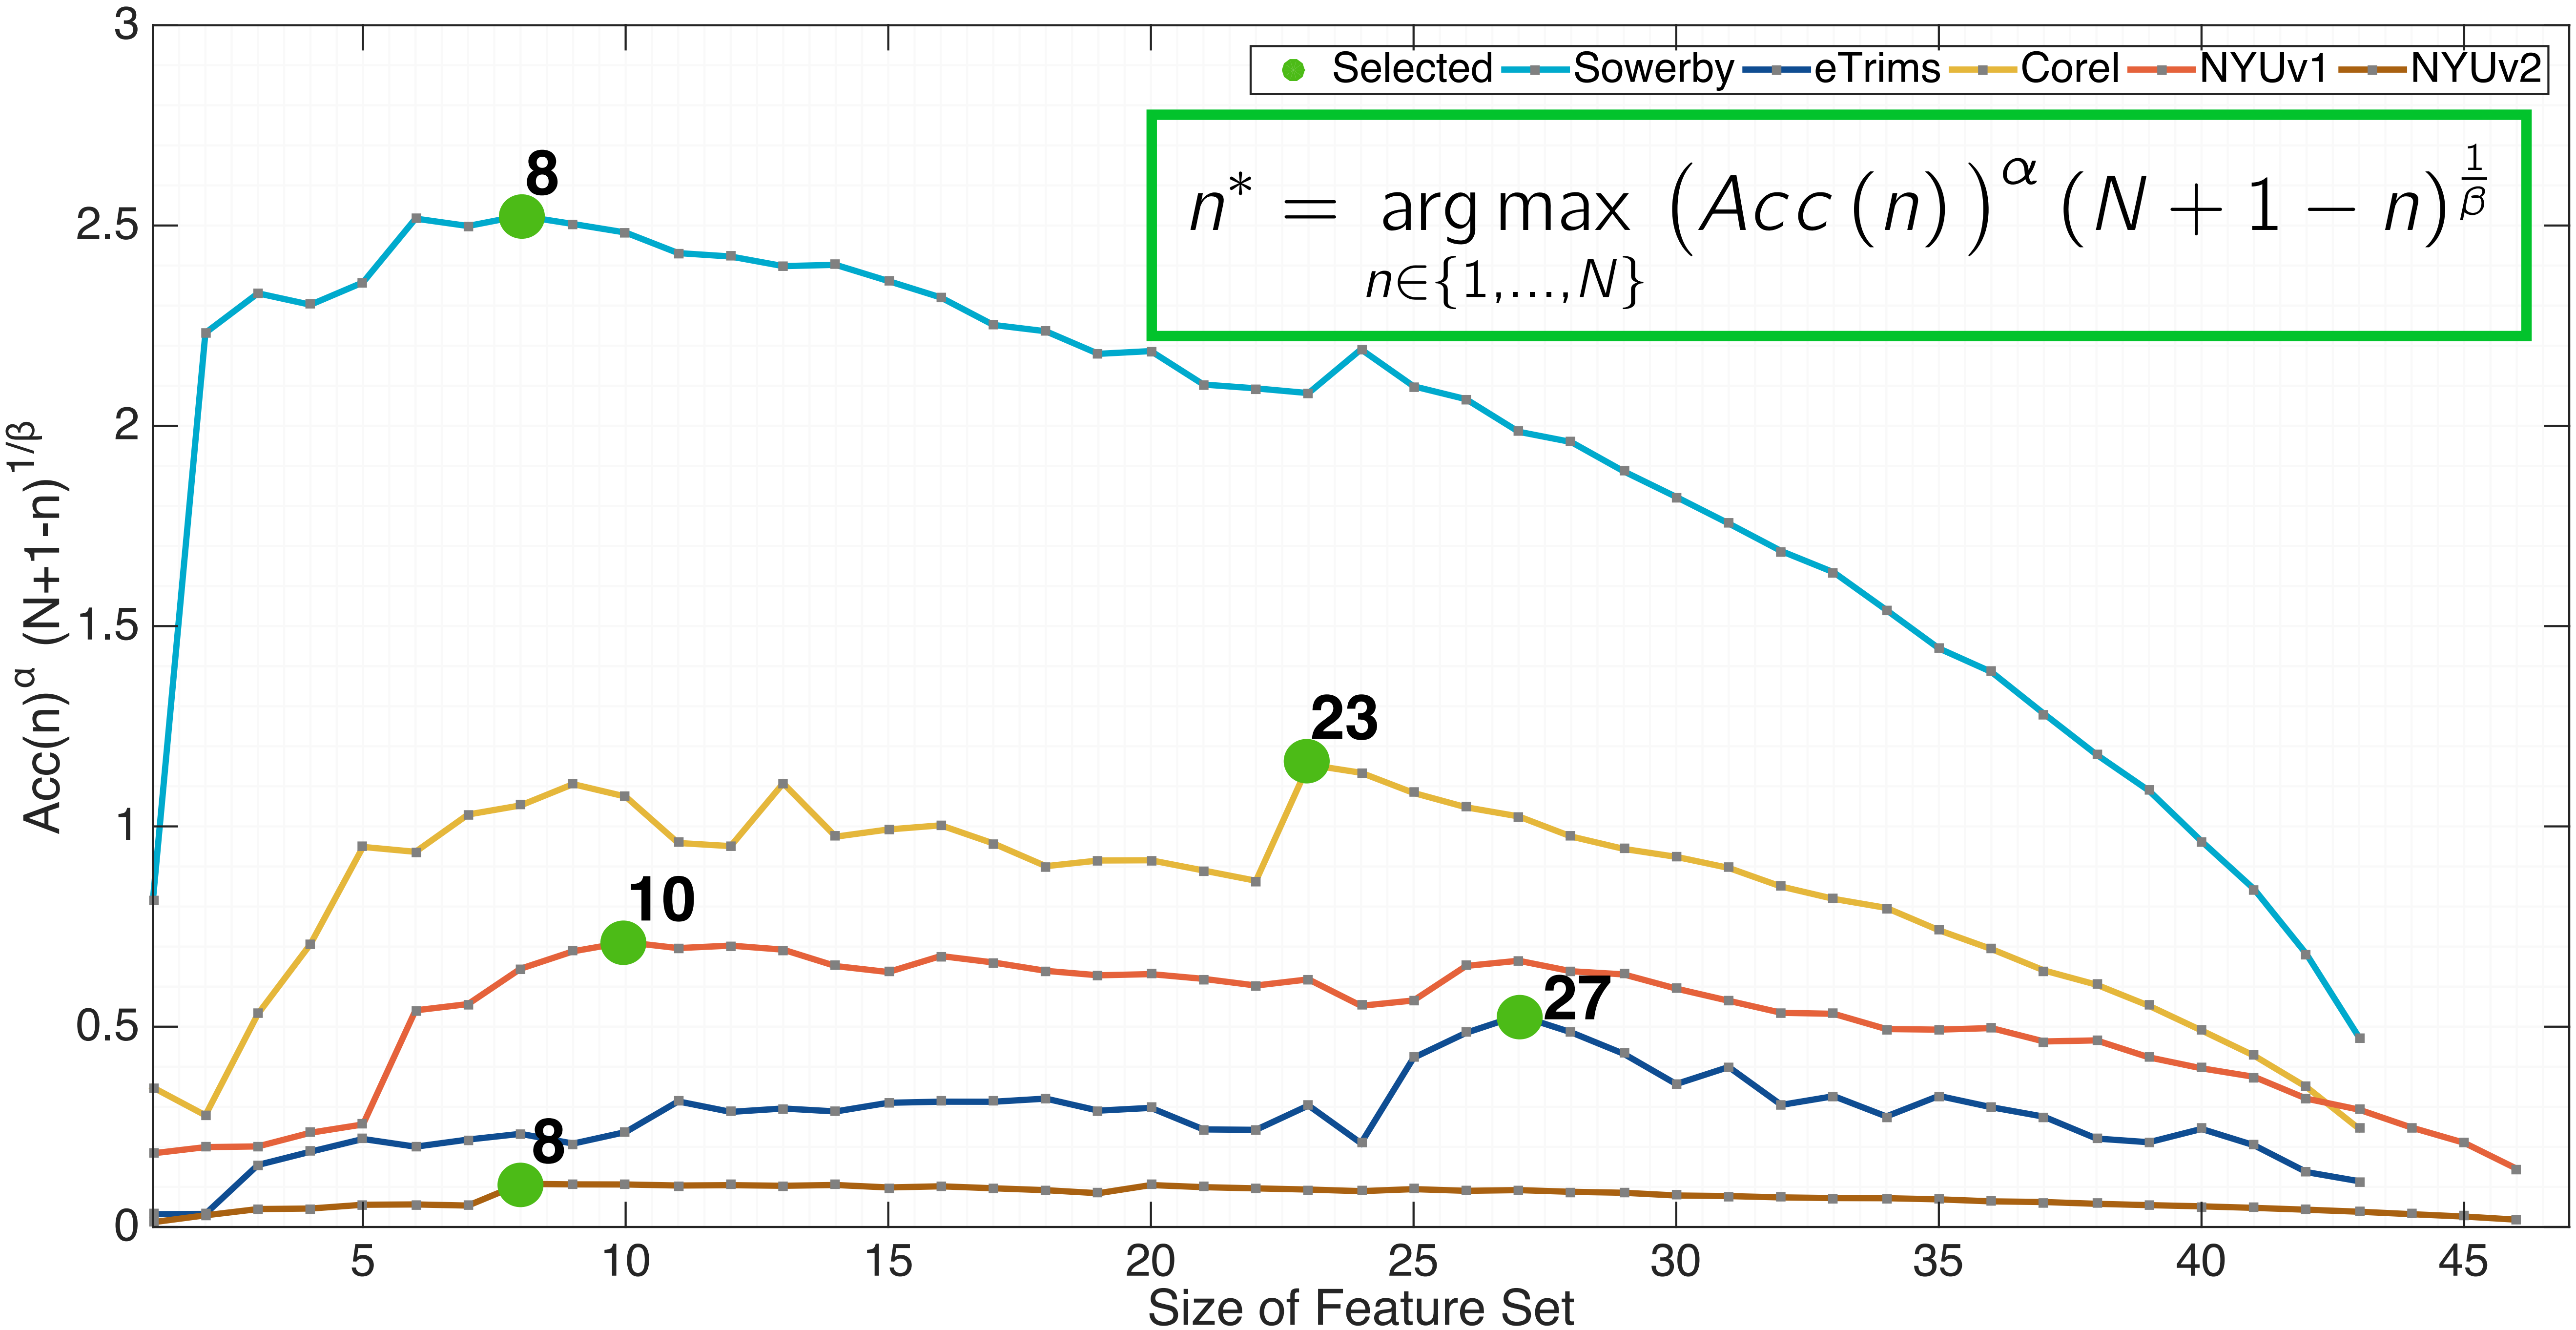

Semantic Image Segmentation

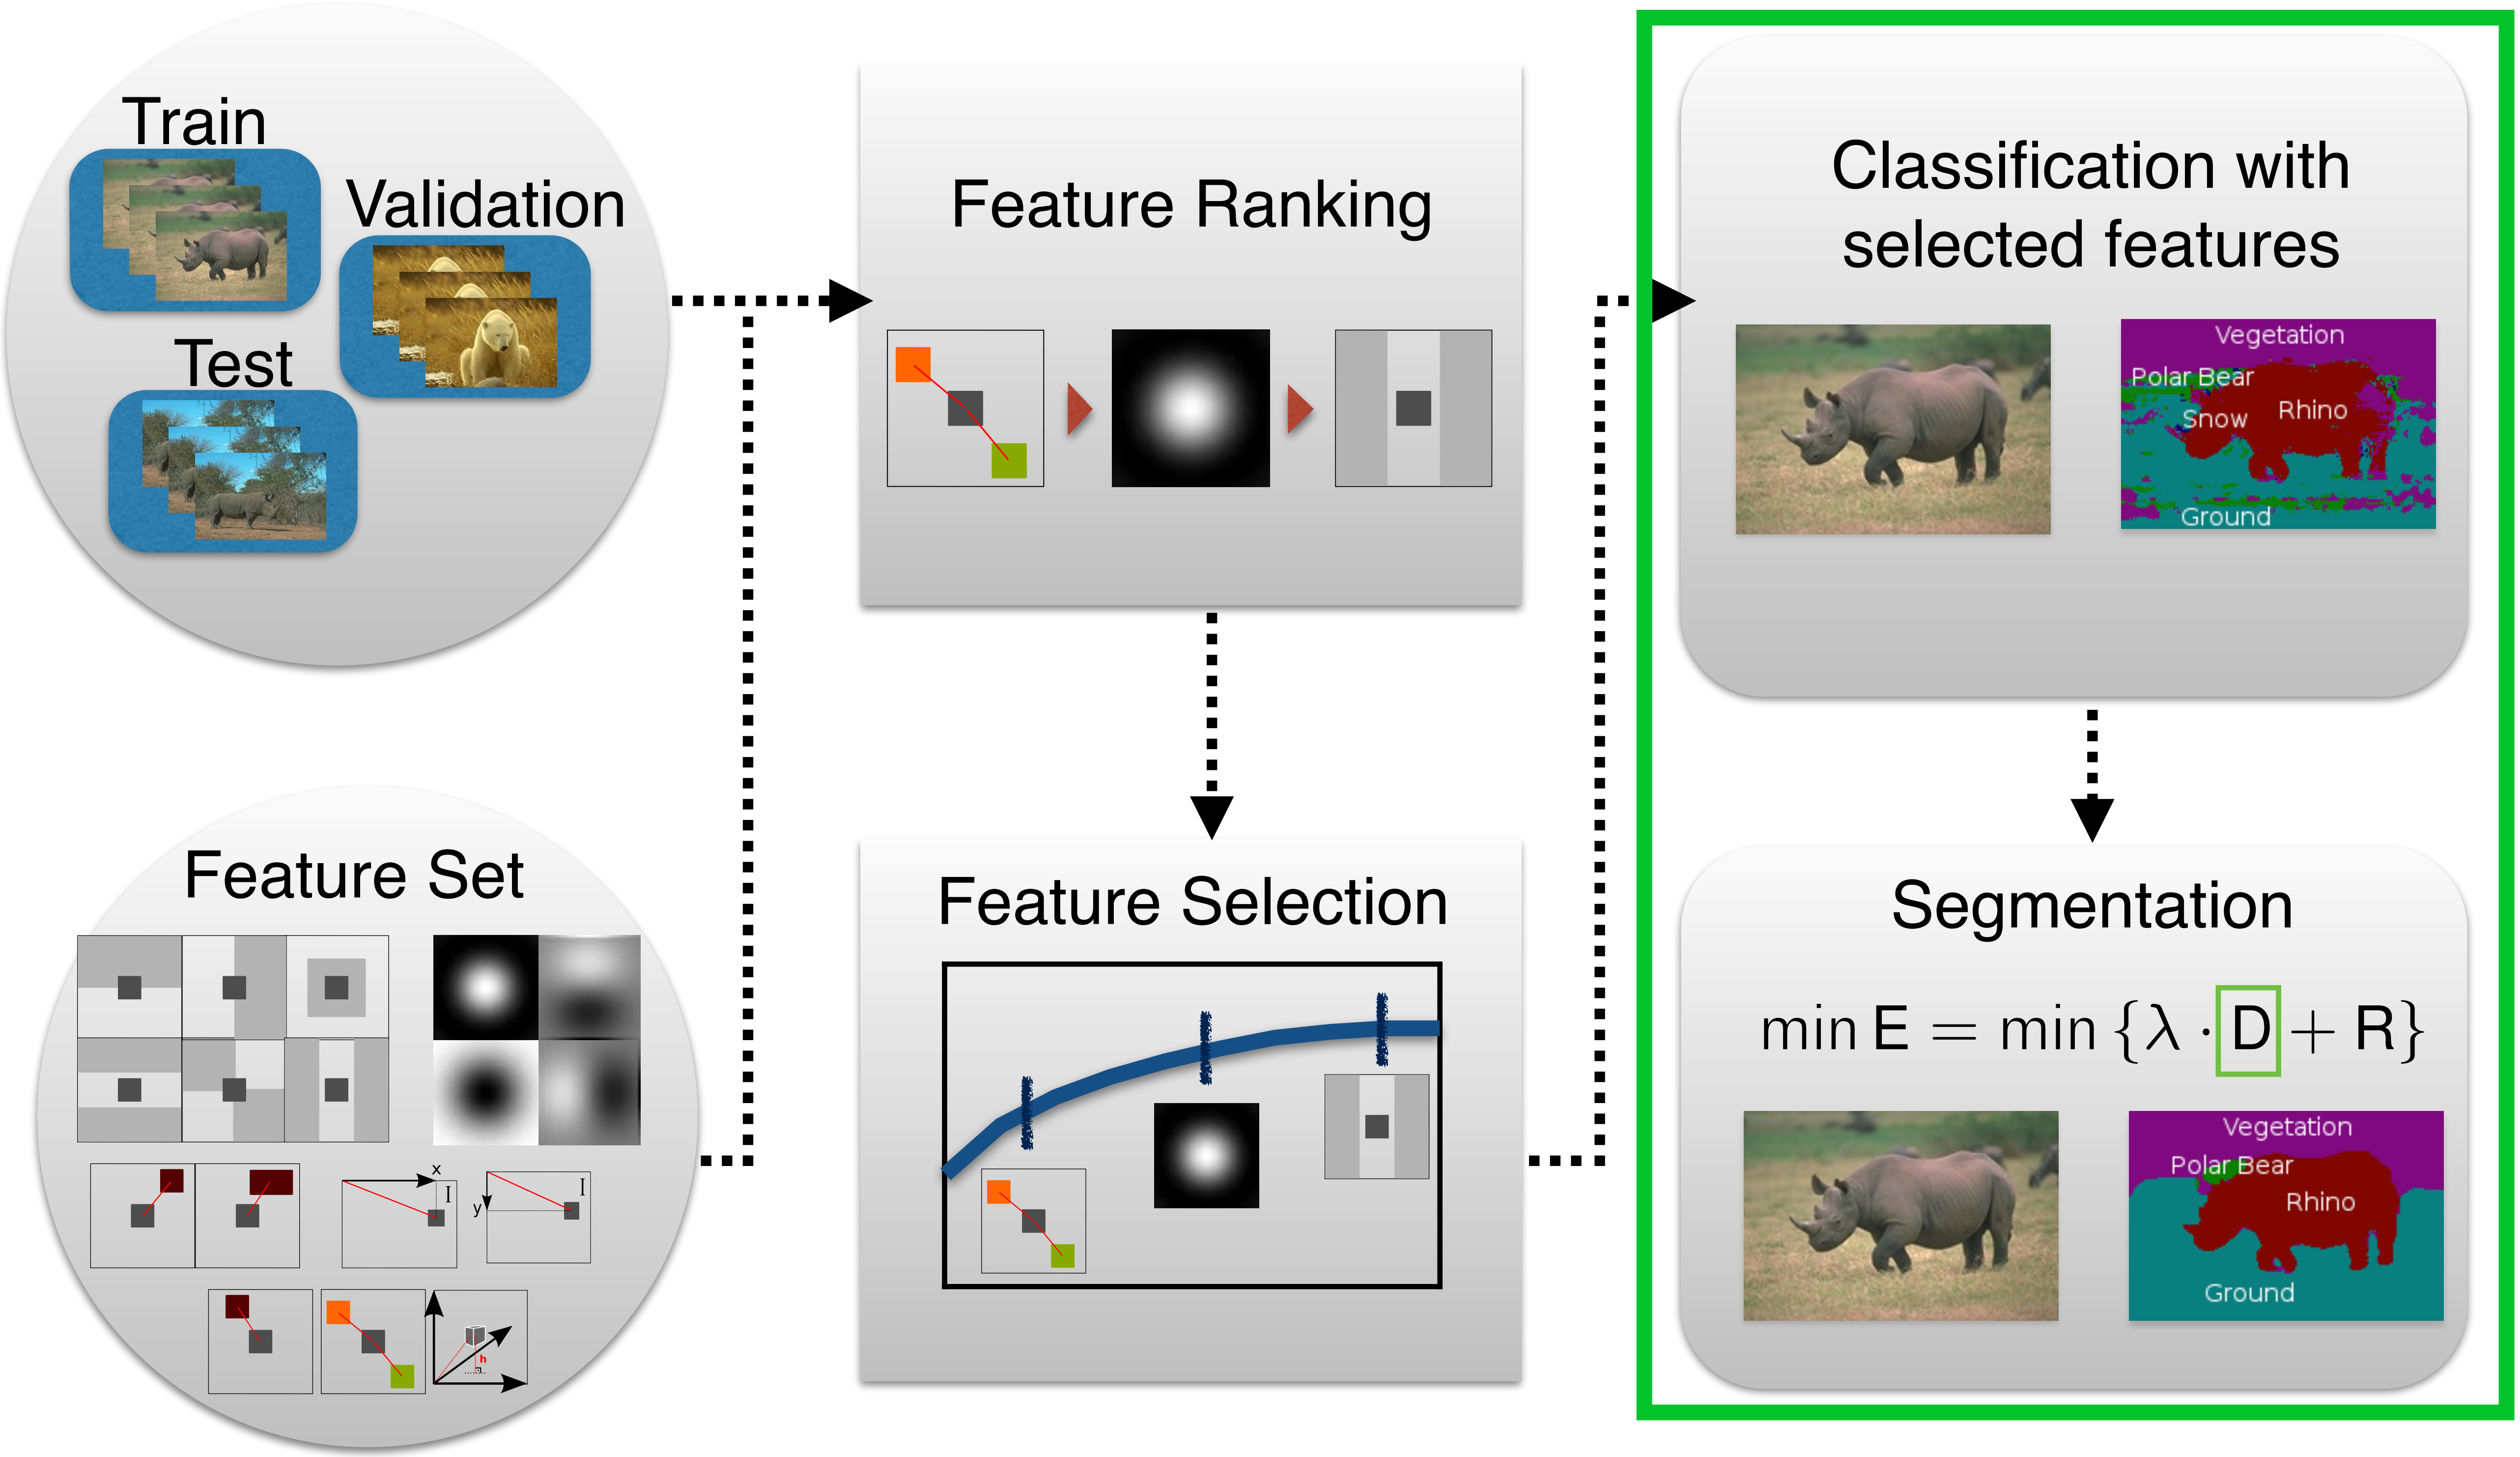

Classification & Segmentation

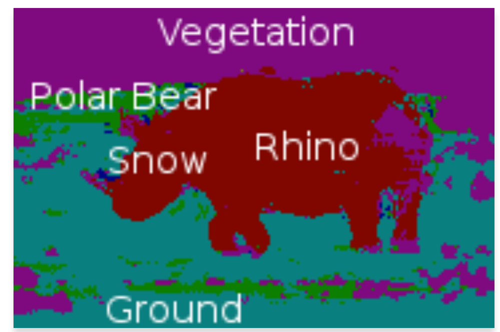

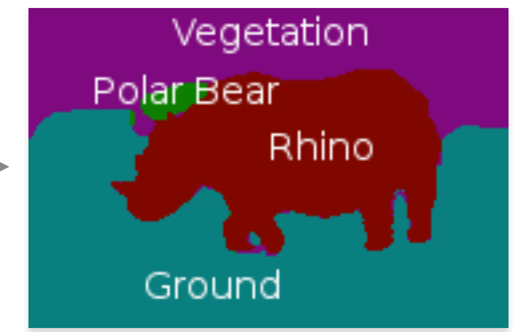

arg max D

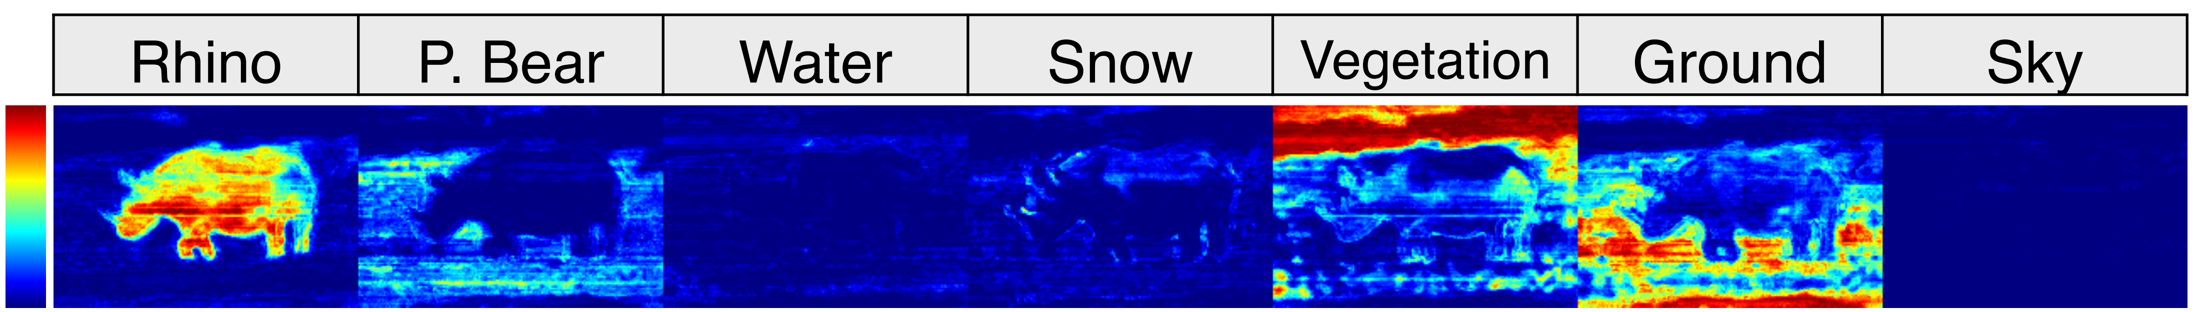

$$\min \mathrm{E} = \min \{\lambda \cdot \boxed{\mathrm{D}} + \mathrm{R}\}$$

Improved Runtime

	Training Time (in seconds)					Testing Time (in seconds)				
	eTrims	Corel	Sowerby	NYUv1	NYUv2	eTrims	Corel	Sowerby	NYUv1	NYUv2
Shotton et al.	—	1800	1200	—	—	—	1.10	2.50	—	—
Fröhlich et al.	—	—	—	—	—	17.0	—	—	—	—
Couprie et al.	—	—	—	—	172800	—	—	—	—	0.70
Hermans et al.	—	—	—	—	—	—	—	—	0.38	0.38
Proposed	143	20	2	133	183	6.6	0.27	0.07	0.32	0.26

On average we improve the runtime by a factor of 7.7

Competitive Results

	Classification					Segmentation				
	eTrims	Corel	Sowerby	NYUv1	NYUv2	eTrims	Corel	Sowerby	NYUv1	NYUv2
Shotton et al.	—	68.4	85.6	—	—	—	74.6	88.6	—	—
Fröhlich et al.	—	—	—	—	—	77.2	—	—	—	—
Couprie et al.	—	—	—	—	—	—	—	—	—	52.4
Hermans et al.	—	—	—	65.0	—	—	—	—	71.5	54.2
Proposed	77.1	74.4	87.1	65.0	44.0	77.9	78.2	88.8	66.5	45.0

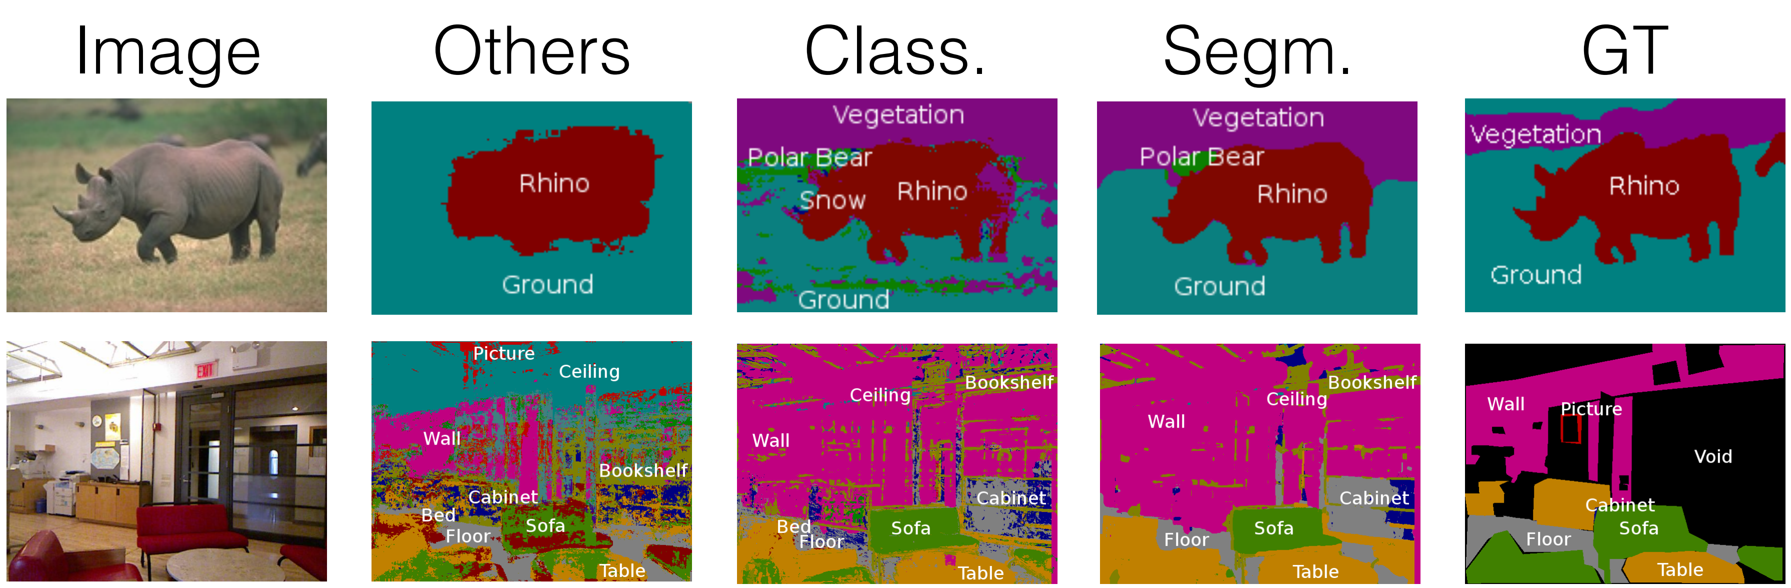

Optimizing the Relevance-Redundancy Tradeoff for Efficient Semantic Segmentation

Caner Hazırbaş | hazirbas@cs.tum.edu

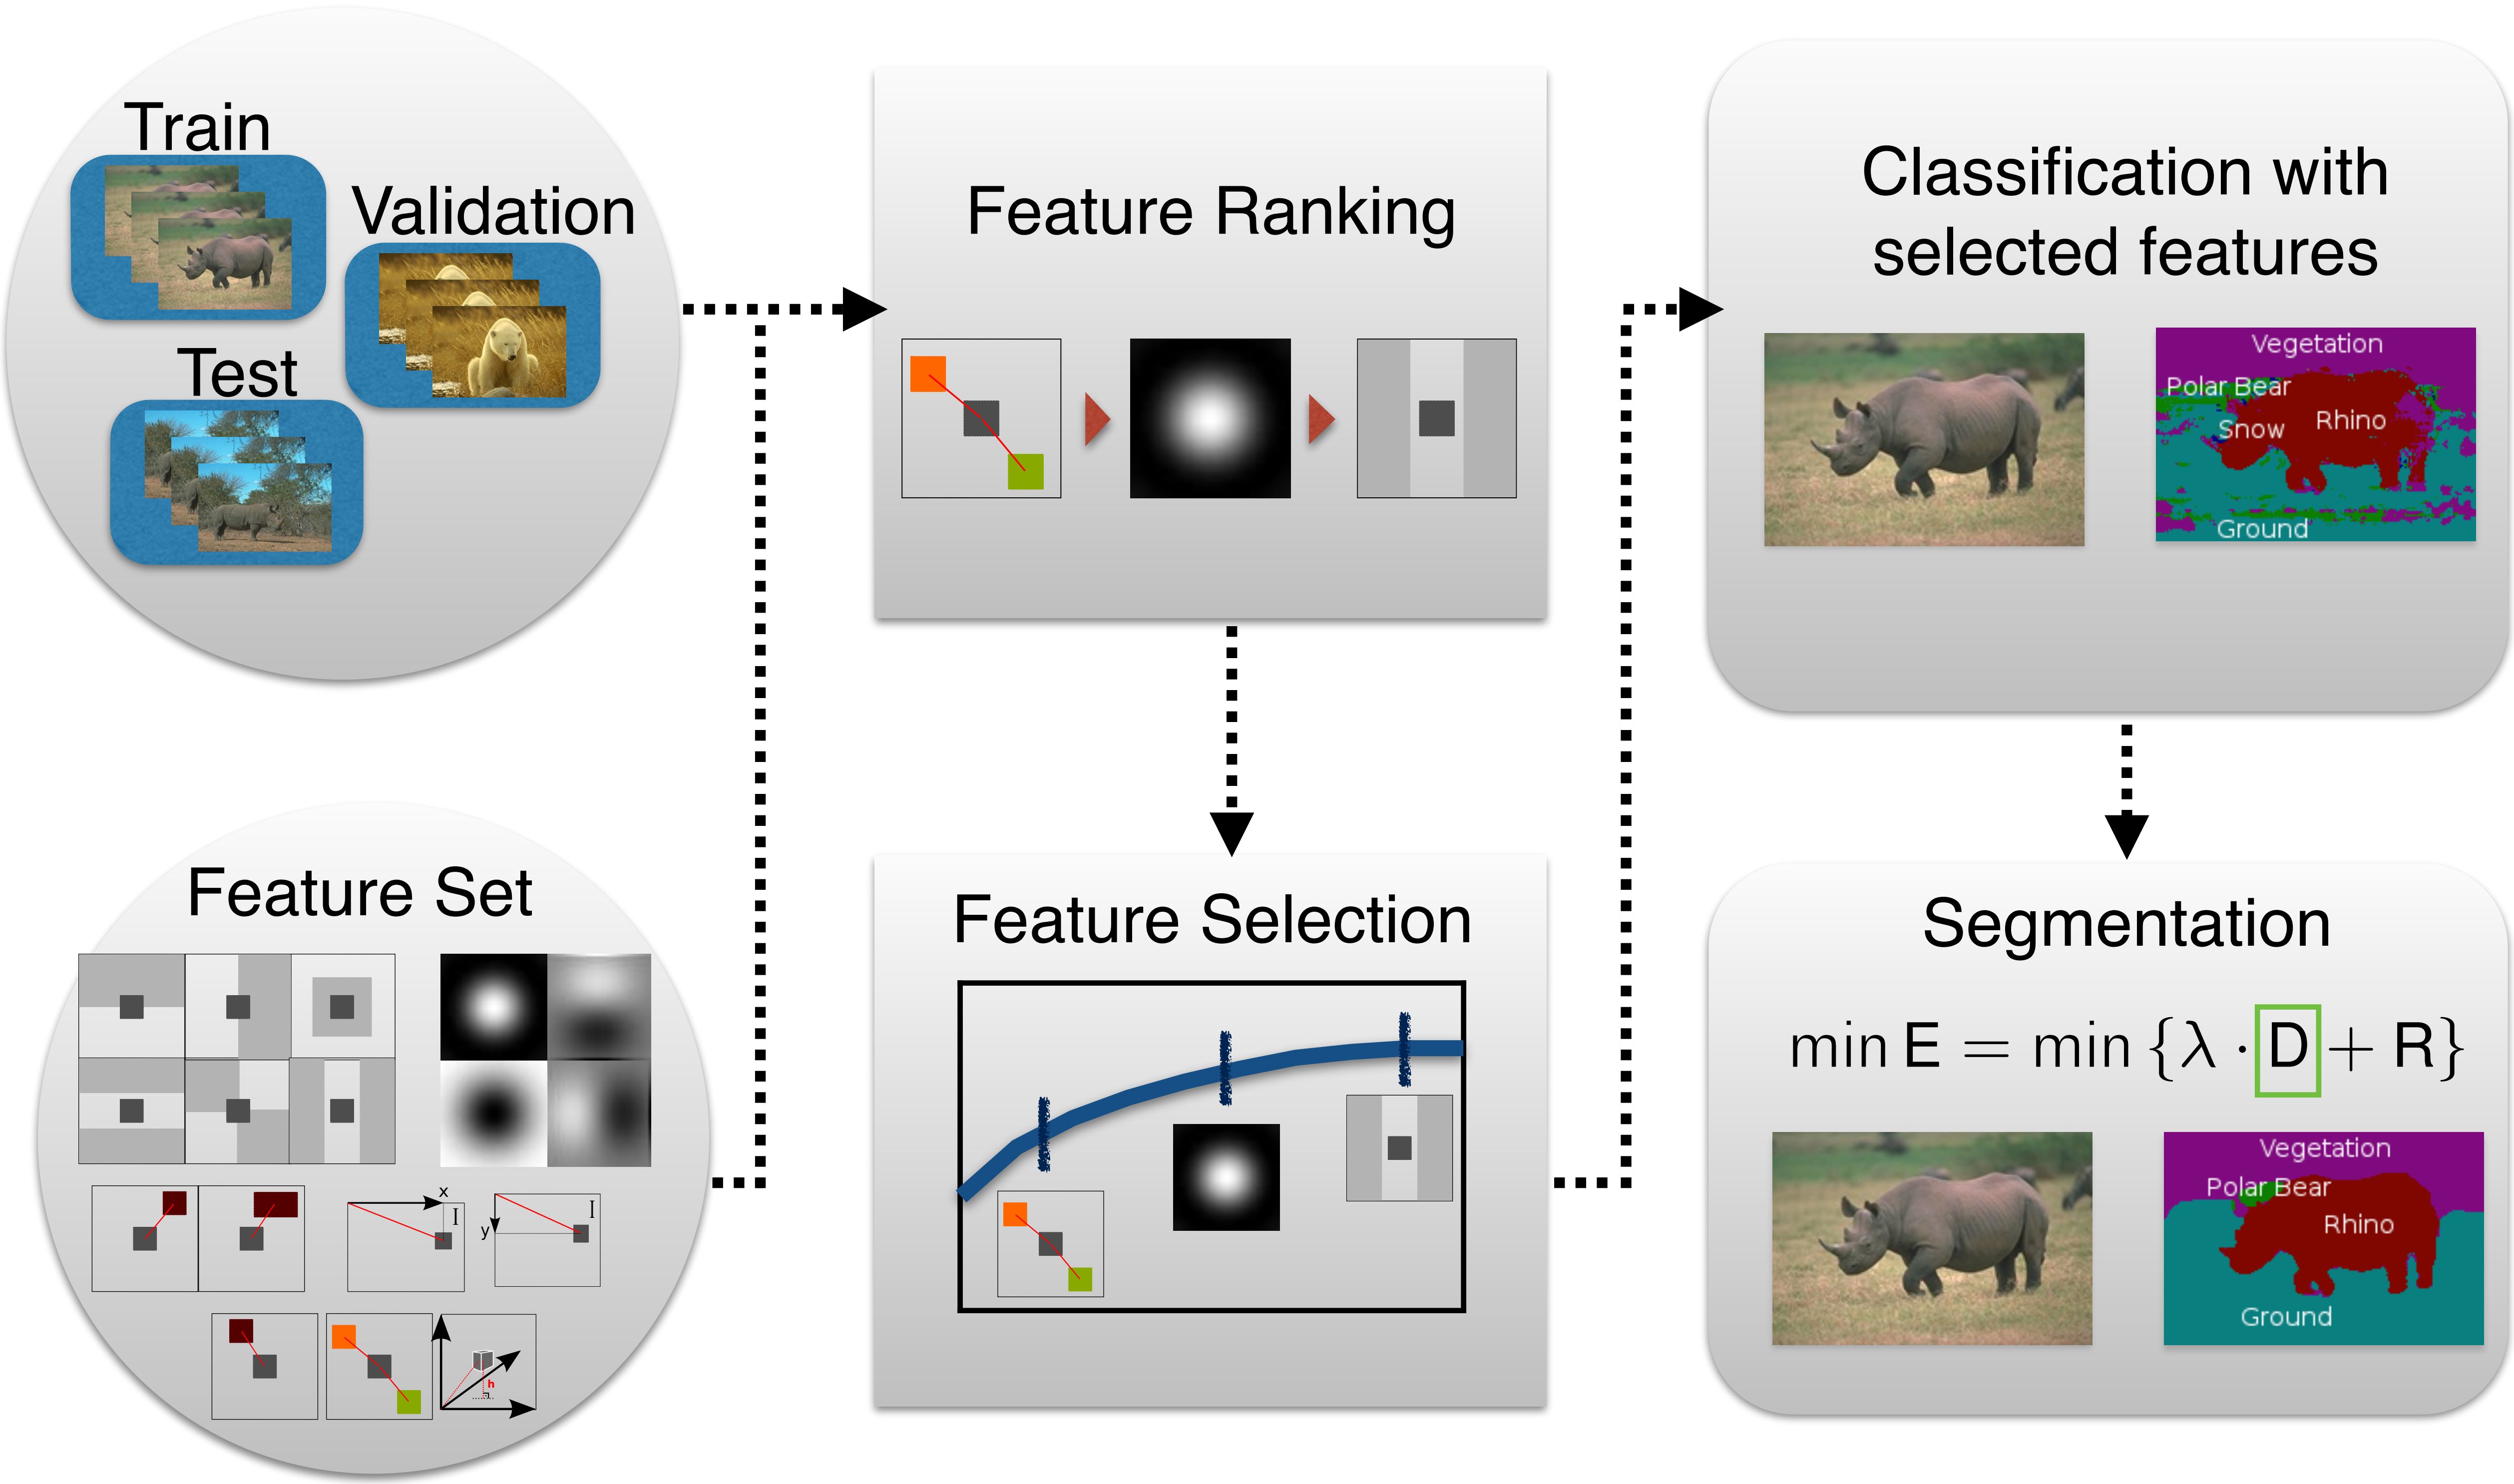

Optimizing the Relevance-Redundancy Tradeoff for Efficient Semantic Segmentation

Caner Hazırbaş | hazirbas@cs.tum.edu

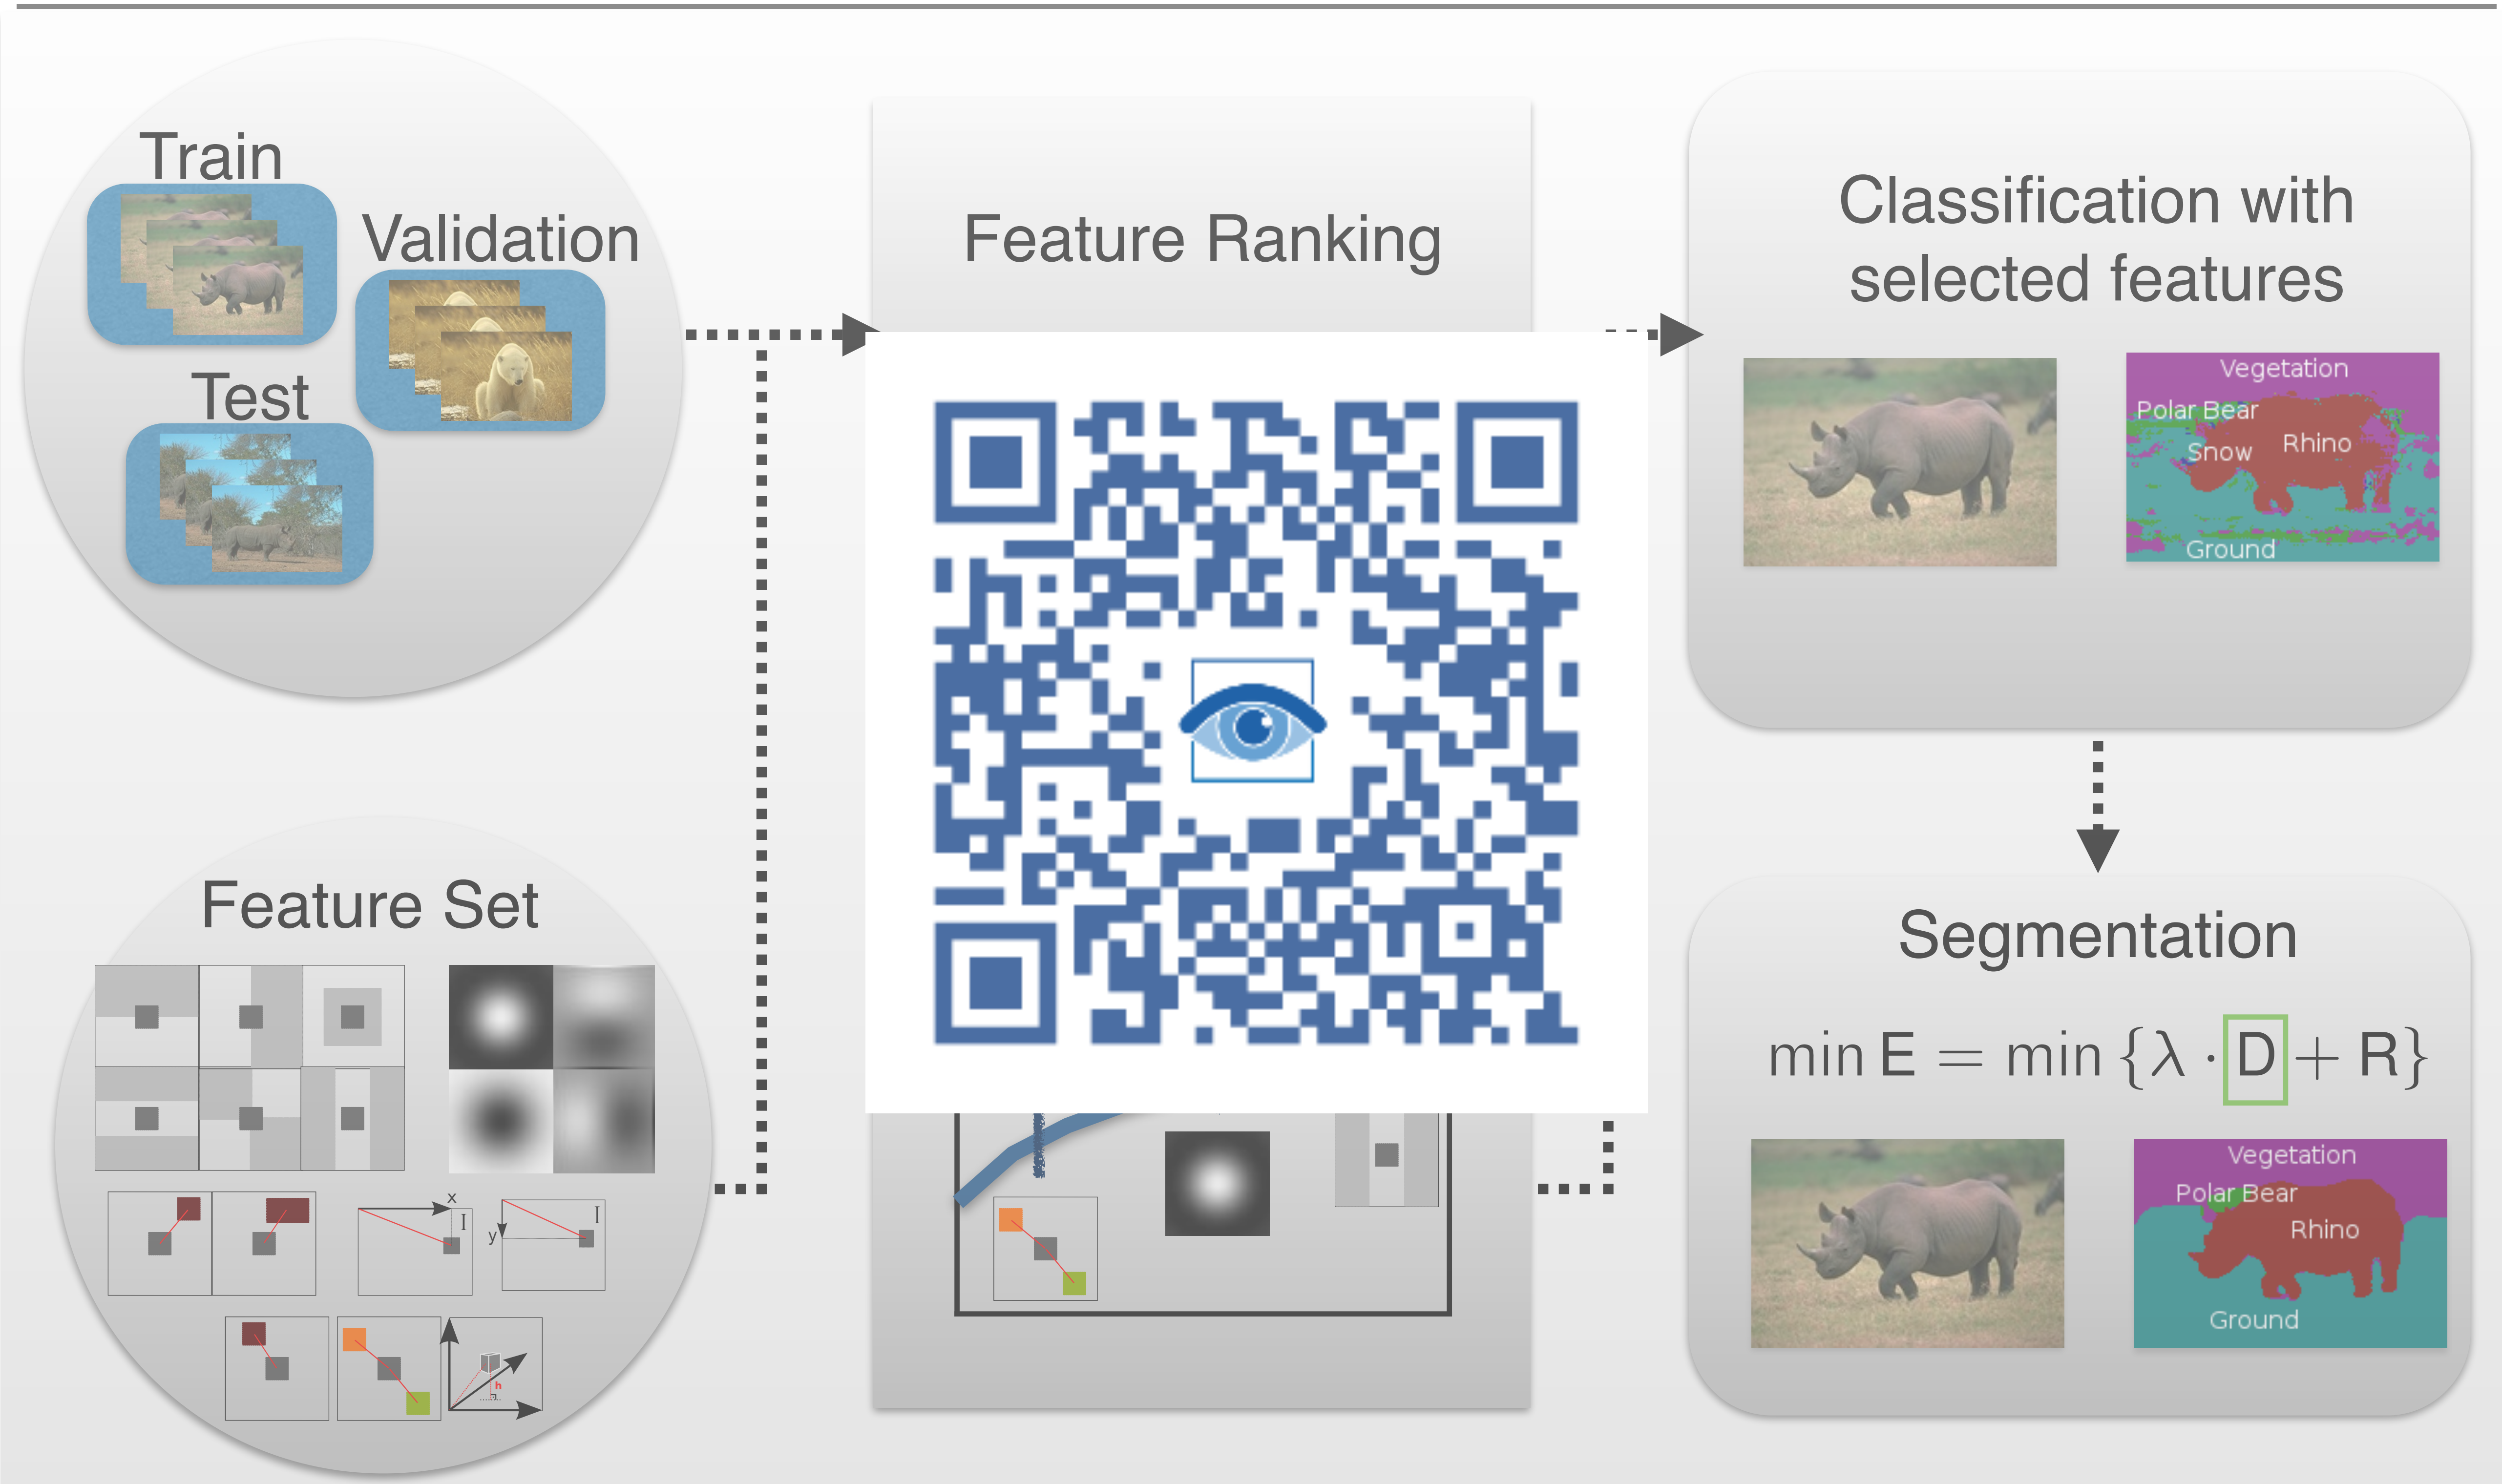

References

- TextonBoost for Image Understanding: Multi-Class Object Recognition and Segmentation by Jointly Modeling Texture, Layout, and Context
  Jamie Shotton, John Winn, Carsten Rother, and Antonio Criminisi IJCV'07
- Semantic Segmentation with Millions of Features: Integrating Multiple Cues in a Combined Random Forest Approach
  Björn Fröhlich, Erik Rodner, and Joachim Denzler ACCV'12
- Indoor Semantic Segmentation using depth information
  Camille Couprie, Clément Farabet, Laurent Najman, and Yann LeCun ICLR'13
- Dense 3D Semantic Mapping of Indoor Scenes from RGB-D Images
  Alexander Hermans, Georgios Floros and Bastian Leibe ICRA'14