

Large-Margin Determinantal Point Processes

Wei-Lun Chao^{*1}, Boqing Gong^{*1}, Kristen Grauman², and Fei Sha¹

¹University of Southern California, ²University of Texas at Austin

Highlights

- Investigate determinantal point processes (DPPs) for discriminative subset selection
- Propose **margin** based parameter estimation to explicitly track errors in selecting subsets
- Balance different types of evaluation metrics, e.g., precision and recall
- Improve modeling flexibility by multiple-kernel based parameterization
- Attain state-of-the-art performance on the tasks of video and document summarization

Background

- A DPP defines a probabilistic distribution over the power set of a ground set: *diverse subsets with large probabilities*

Ground set of M items, $\mathcal{Y} = \{1, 2, \dots, M\}$

$\mathbf{L} \in \mathbb{S}_+^M$: a kernel matrix of pairwise similarities

$$P(\mathbf{y} \subseteq \mathcal{Y}; \mathbf{L}) = \frac{\det(\mathbf{L}_{\mathbf{y}})}{\det(\mathbf{L} + \mathbf{I})}$$

$$P(\mathbf{y} = \{i, j\}; \mathbf{L}) \propto \det(\mathbf{L}_{\{i, j\}}) = L_{ii}L_{jj} - L_{ij}^2$$

- DPPs offer a powerful approach to modeling **diversity** in applications where the goal is to select a diverse subset from a ground set of items (e.g., retrieval, summarization)
- MAP inference (NP-hard): $\mathbf{y}^{\text{MAP}} = \operatorname{argmax}_{\mathbf{y}} P(\mathbf{y}; \mathbf{L})$
- Estimate the kernel \mathbf{L} from labeled data $\{(\mathbf{y}^{*(n)}, \mathcal{Y}^{(n)})\}$
 - Reparameterization: $\mathbf{L}^{(n)}(\mathcal{Y}^{(n)}; \cdot)$
 - Standard method: Maximum likelihood estimation (MLE)

$$\mathbf{L}^{\text{MLE}} = \operatorname{argmax}_{\mathbf{L}} \sum_n \log P(\mathbf{y}^{*(n)}; \mathbf{L}^{(n)}(\mathcal{Y}^{(n)}; \cdot))$$

Problems with existing methods

Statistical challenges

- Limited number of training samples

Modelling challenges

- Limited power in parameterizing kernels with the widely used quality-diversity (QD) decomposition
- Unable to track discriminative errors in selecting subsets
- Unable to differentiate different types of metrics (e.g., precision, recall) for complex structured prediction tasks

Contribution I: Multiple-kernel representation (MKR)

- Quality-diversity (QD) decomposition

$$\forall i \in \mathcal{Y} \left\{ \begin{array}{l} \mathbf{x}_i : \text{quality features} \\ \mathbf{w}_i : \text{similarity features} \end{array} \right. \Rightarrow \left\{ \begin{array}{l} L_{ij} = q_i q_j S_{ij} = q_i q_j \mathbf{w}_i^T \mathbf{w}_j \\ q_i = q(\mathbf{x}_i) = \exp(\mathbf{w}_i^T \mathbf{x}_i) \end{array} \right.$$

- Multiple-kernel representation (MKR)

$$S_{ij} = \sum_k r_k \exp\left\{-\|\mathbf{w}_i - \mathbf{w}_j\|_2^2 / \tau_k^2\right\} + s \mathbf{w}_i^T \mathbf{w}_j, \quad \text{s.t. } \sum_k r_k + s = 1$$

Selected references:

- [1] A. Kulesza and B. Taskar. Determinantal point processes for machine learning. 2012.
- [2] H. T. Dang. Overview of DUC 2005. In Document Understanding Conf., 2005.
- [3] S. E. F. de Avila et al. VSUMM: A mechanism designed to produce static video summaries and a novel evaluation method. Pattern Recognition Letters, 2011.

Contribution II: Margin based parameter estimation (LME)

- Maintain desired margins between correct/incorrect subsets

$$\log P(\mathbf{y}^*; \mathbf{L}) \geq \max_{\mathbf{y} \subseteq \mathcal{Y}} \log \{\ell(\mathbf{y}^*, \mathbf{y}) P(\mathbf{y}; \mathbf{L})\}$$

$$= \max_{\mathbf{y} \subseteq \mathcal{Y}} \log \ell(\mathbf{y}^*, \mathbf{y}) + \log P(\mathbf{y}; \mathbf{L})$$

➤ The multiplicative margin leads to tractable optimization
- Measure the subset discrepancy by structured loss functions

$$\ell_S(\mathbf{y}^*, \mathbf{y}) = \underbrace{\sum_{i \notin \mathbf{y}^*} \mathbb{I}[i \in \mathbf{y}]}_{\text{precision}} + \underbrace{\sum_{i \in \mathbf{y}^*} \mathbb{I}[i \notin \mathbf{y}]}_{\text{recall}}$$
- Optimization: Jensen's inequality (softmax) for tractability

$$\log P(\mathbf{y}^*; \mathbf{L}) \geq \operatorname{softmax}_{\mathbf{y} \subseteq \mathcal{Y}} \log \ell_S(\mathbf{y}^*, \mathbf{y}) + \log P(\mathbf{y}; \mathbf{L})$$

$$= \log \left(\sum_{i \notin \mathbf{y}^*} K_{ii} + \sum_{i \in \mathbf{y}^*} (1 - K_{ii}) \right), \text{ where } K = \mathbf{L}(\mathbf{L} + \mathbf{I})^{-1}$$
- Objective function: hinge loss $[\cdot]_+ = \max(0, \cdot)$

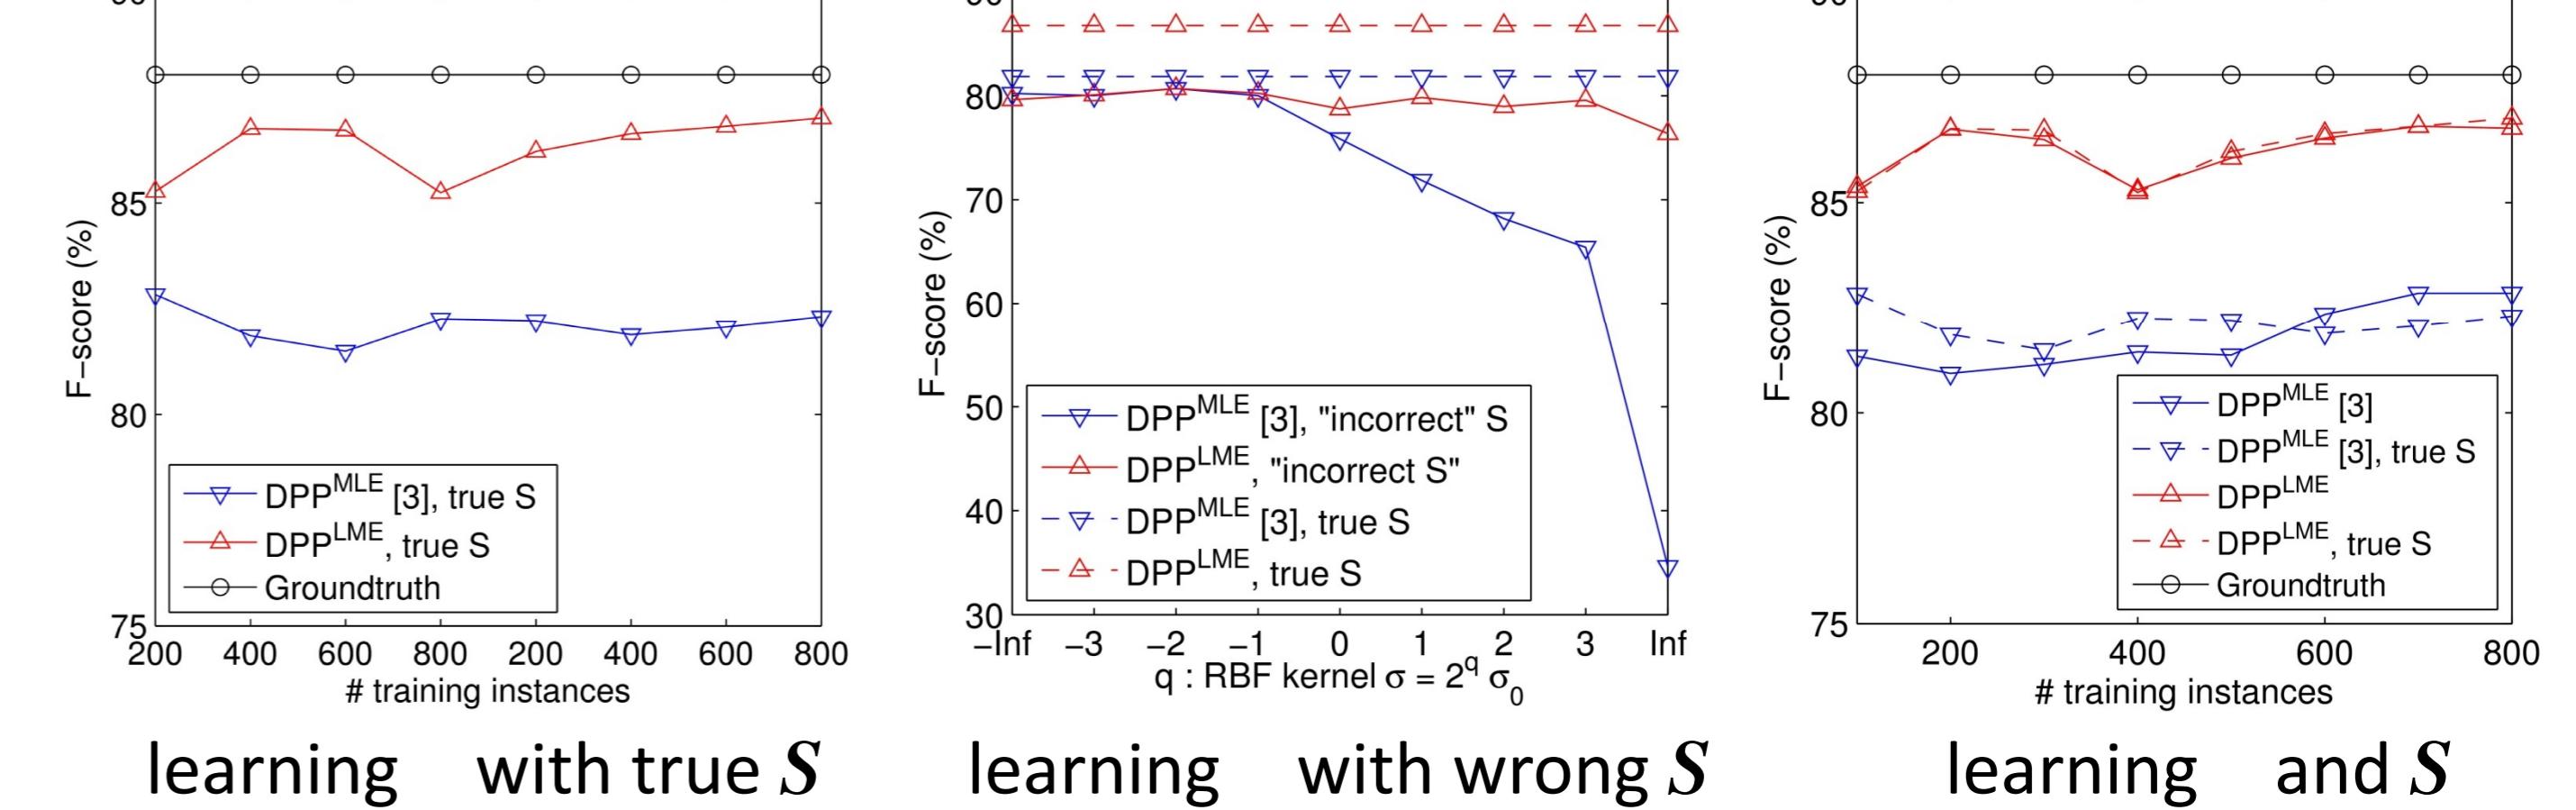
$$\min \sum_n \left[-\log P(\mathbf{y}^{*(n)}; \mathbf{L}^{(n)}) + \log \left(\sum_{i \notin \mathbf{y}^{*(n)}} K_{ii}^{(n)} + \sum_{i \in \mathbf{y}^{*(n)}} (1 - K_{ii}^{(n)}) \right) \right]_+$$

Experiments

Evaluation: Precision, Recall, F-score (harmonic mean of P, R)

Inference: Brute-force search, minimum Bayesian risk (MBR)

Synthetic dataset:

- Generate ground sets and target subsets by sampling $\{\mathbf{x}_i = \mathbf{z}_i\}$, computing $S_{ij} = q_i q_j$, brute-force search, and adding noise
- 
- Our method (**MKR+LME**) is more robust in parameter estimation

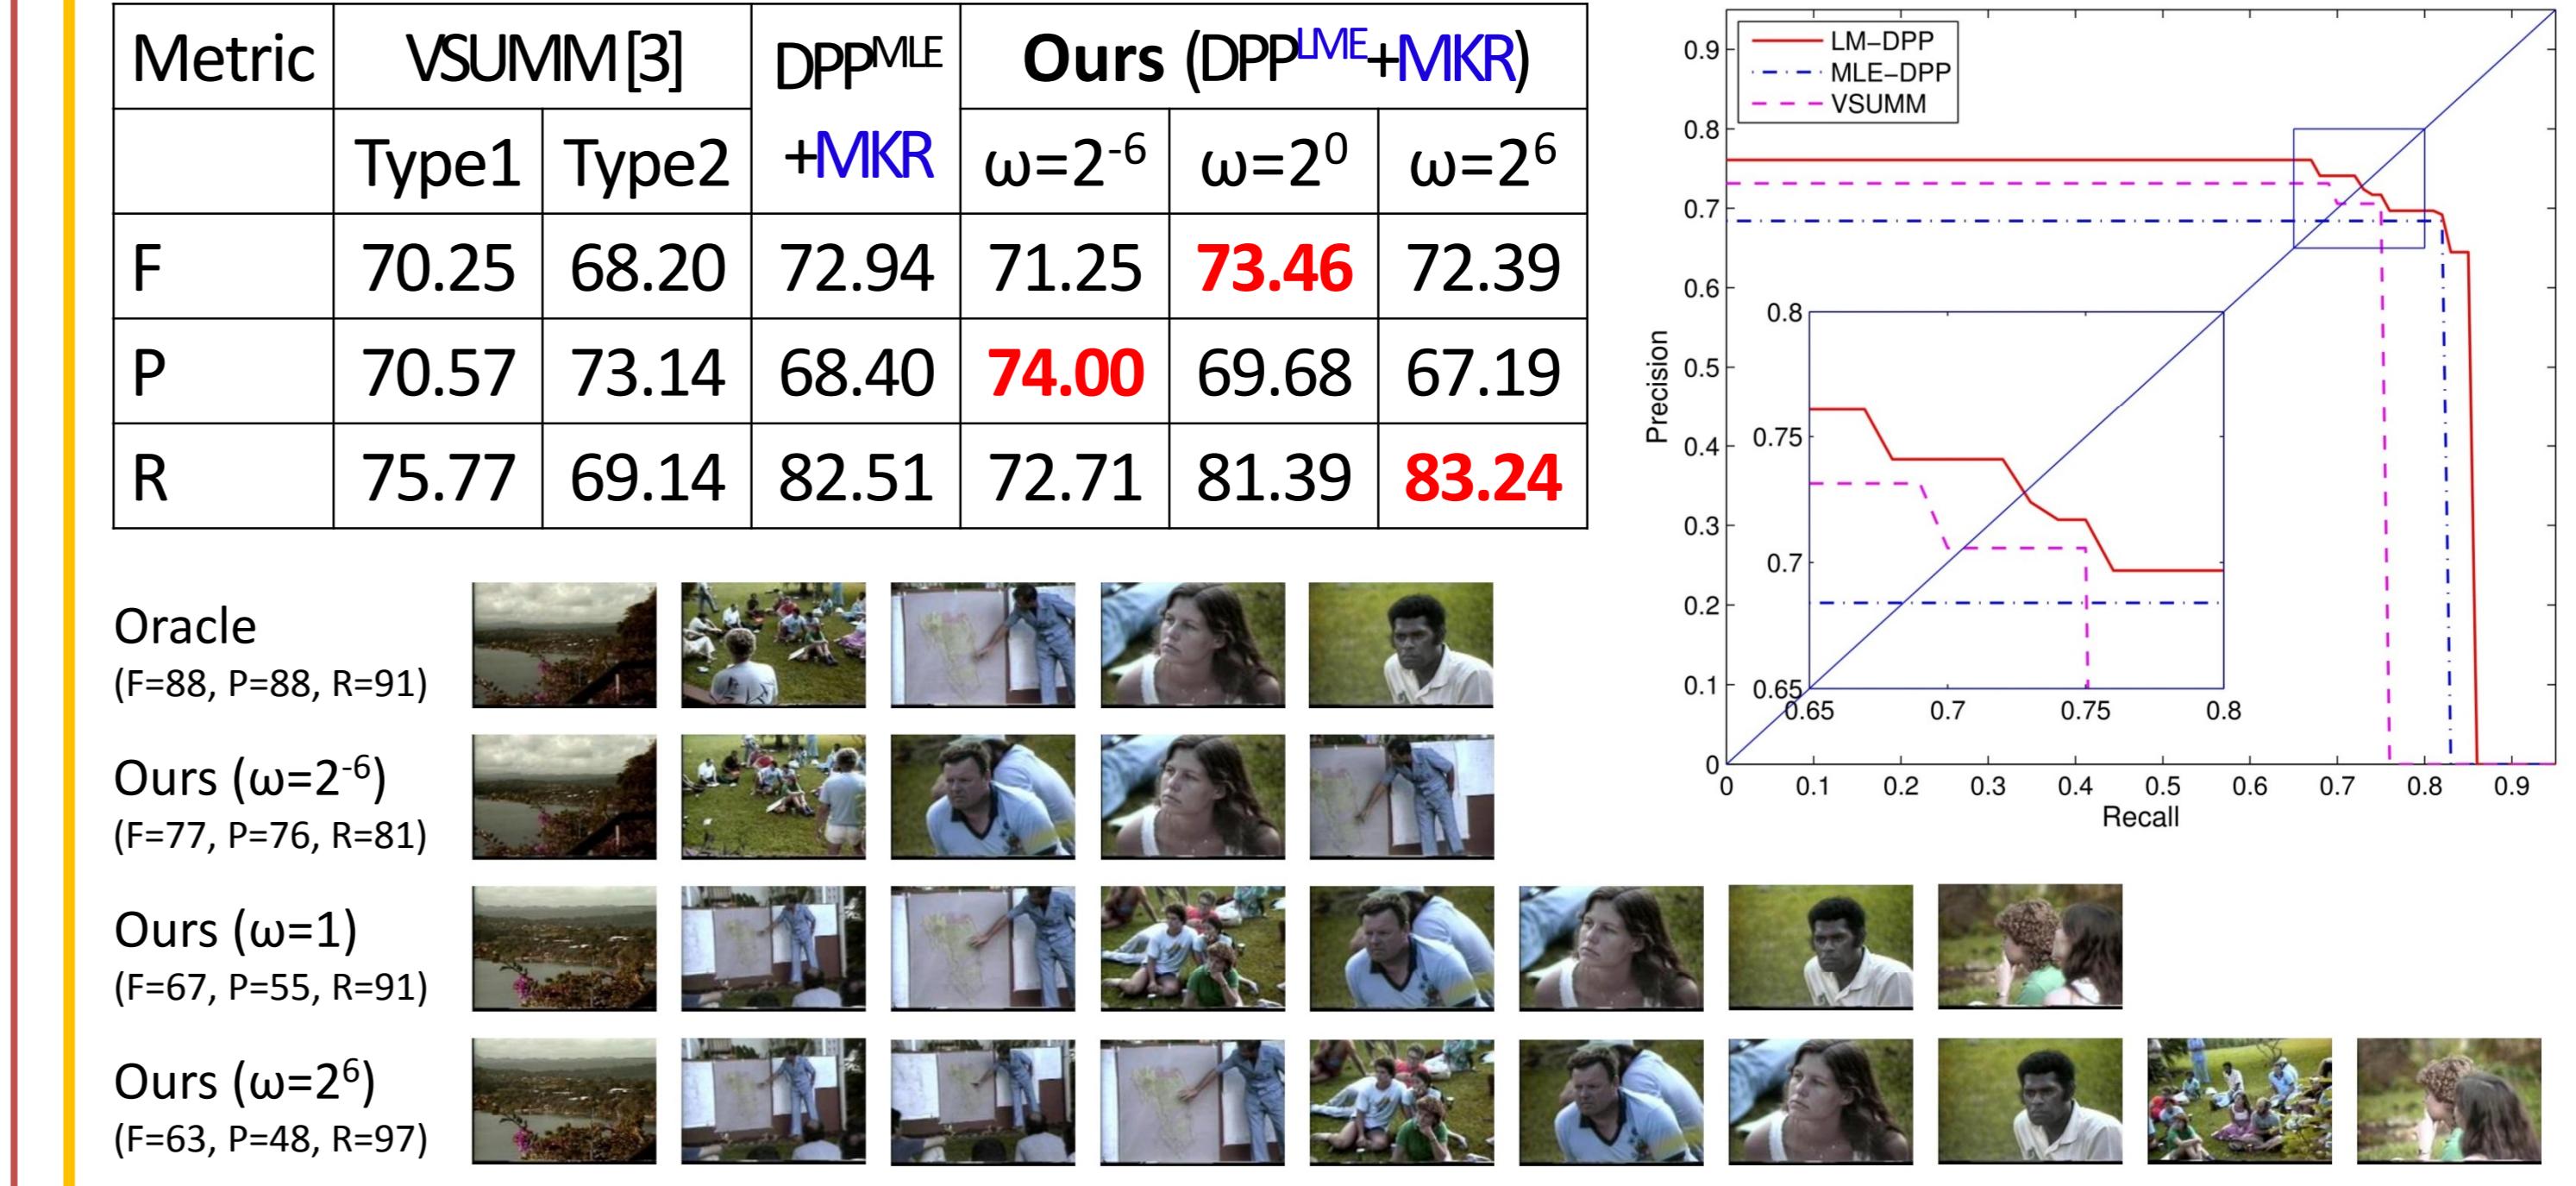
Document summarization: Document Understanding Conf. (DUC)

Train: DUC 2003 (60 clusters) Test: DUC 2004 (50 clusters)

Method	ROUGE1-F	ROUGE1-P	ROUGE1-R	ROUGE2-F	ROUGE2-P	ROUGE2-R
PEER65 [2]	37.9	37.6	38.2	9.13	--	--
DPP ^{MLE} +COS	37.9±0.08	37.4±0.08	38.5±0.08	7.72±0.06	7.63±0.06	7.83±0.06
Ours (DPP ^{LME} +COS)	38.4±0.09	37.7±0.10	39.1±0.08	8.20±0.07	8.07±0.07	8.35±0.07
Ours (DPP ^{MLE} +MKR)	39.1±0.08	39.0±0.09	39.3±0.09	9.25±0.08	9.24±0.08	9.27±0.08
Ours (DPP ^{LME} +MKR)	39.7±0.05	39.6±0.08	39.9±0.06	9.40±0.08	9.38±0.08	9.43±0.08

Video summarization: Open Video Project (OVP)

50 videos from OVP, 5 user annotations, 5-fold cross-validation



Our method can balance different types of evaluation metrics and achieve better performance in both summarization tasks