

This is a repository copy of *Quantum vs classical ranking in segment grouping*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/85364/>

Version: Accepted Version

Proceedings Paper:

Escolano, Francisco, Bonev, Boyan and Hancock, Edwin R. orcid.org/0000-0003-4496-2028 (2014) Quantum vs classical ranking in segment grouping. In: Structural, Syntactic, and Statistical Pattern Recognition: Joint IAPR International Workshop, S+SSPR 2014, Joensuu, Finland, August 20-22, 2014. Proceedings. Joint IAPR International Workshop on Structural, Syntactic, and Statistical Pattern Recognition, S+SSPR 2014, 20-22 Aug 2014 Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) . Springer , GBR , pp. 203-212.

https://doi.org/10.1007/978-3-662-44415-3_21

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Quantum vs Classical Ranking in Segment Grouping

Francisco Escolano, Boyan Bonev, and Edwin R. Hancock

Department of Computer Science and AI, University of Alicante, Spain

Department of Statistics, UCLA, USA

Department of Computer Science, University of York, UK

Abstract. In this paper we explore the use of ranking as a mean of guiding unsupervised image segmentation. Starting by the well known Pagerank algorithm we introduce an extension based on quantum walks. Pagerank (rank) can be used to prioritize the merging of segments embedded in uniform regions (parts of the image with roughly similar appearance statistics). Quantum Pagerank, on the other hand, gives high priority to boundary segments. This latter effect is due to the higher order interactions captured by quantum fluctuations. However we found that qrank does not always outperform its classical version. We analyze the Pascal VOC database and give Intersection over Union (IoU) performances.

Keywords: random walks, quantum walks, ranking, segment grouping

1 Introduction

When applied to image segmentation, random walks have been used to propagate labels in a semi-supervised way. For instance, in [1] pixels are labeled in terms of the probability that a random walk will reach them from a given seed.

However, the random walker approach assumes that the weighting function quantifying the dissimilarity between pixel intensities is symmetric. From a graph theoretic perspective this simplifies the problem since the Laplacian matrices of undirected graphs are semi-definite positive. At the same time the resulting asymmetric dissimilarity functions are richer since their directionality allows us to deal with special cases which are particularly interesting in image segmentation. For instance, a symmetric dissimilarity between adjacent segments (e.g. superpixels) imposes a misleading transitivity which may lead to an incorrect grouping. In Fig. 1 (bottom-left), segment X is very compatible with A and B in terms of having similar statistics. However A is in turn more compatible with B than with X. Therefore, A's best candidate for a merging will be B instead of X. This situation also occurs with second-order neighbors (compare X with E).

Incorporating asymmetry into semi-supervised labeling has been done in the area of machine learning. For instance, in [2] conditional probabilities are introduced in the Markov chain, whereas in [3] the graph Laplacian is symmetrized for encoding the directness of the edges. However, to the best of our knowledge

there have been no attempts in the literature to investigate the propagation of information through the digraphs induced by asymmetric dissimilarity measures in an unsupervised context.

In this paper, we explore the impact of both classical and quantum ranking in the selection of the segments to merge in unsupervised segmentation. Our hypothesis is that ranking may improve significantly the quality of the segmentation, since the result of the process contains the information of random or quantum walks probing the network given by the adjacency graph. Therefore, ranking provides local-to-global information that may be critical in a greedy merging process.

The remainder of the paper is organized as follows. In Section 2 we describe the simple hierarchical grouping algorithm used for the study of ranking effects. In Section 3 we review Pagerank from a perspective of digraphs. Section 4 is devoted to the description of the quantum extension of the Pagerank method. Experiments and analysis are presented in Section 5. Finally, in Section 6, we present our conclusions and future work.




Fig. 1. Symmetric vs asymmetric dissimilarity. Top-left: Output of SLIC algorithm. Bottom-left: Asymmetric dissimilarities (in red and blue) between superpixel X and some of its 1^{st} and 2^{nd} -order neighbors (see text). Top-row: Best unsupervised segmentation results imposing symmetry for some VOC Pascal objects. Bottom-row: results by imposing asymmetry. See the scenes analyzed in Top-left of Fig. 2 and Fig. 3.

2 Hierarchical Grouping Algorithm

The algorithm starts with a basic set of segments which are output by the SLIC algorithm [4]. We then build a segmentation hierarchy by merging/composing segments as follows.

Each segment is described by an *appearance vector* $V = (\bar{\mu}, \bar{\sigma}, c_x, c_y, w, h)$, where $\bar{\mu}, \bar{\sigma}$ are the mean and the standard deviation of $(l, a, b, \nabla_x, \nabla_y, \nabla_x^2, \nabla_y^2)$, and (c_x, c_y, w, h) are the centroid of the segment and the dimensions of its bounding box. These appearance vectors are designed so that they can be efficiently computed recursively for new segments composed by merging existing ones.

Each segment has a neighborhood structure. This consists of 1^{st} -order neighbors, which are directly adjacent to the segment, and 2^{nd} -order neighbors (i.e. those adjacent to the 1^{st} -order neighbors). Then, we define an asymmetric similarity function $\Delta_{i|j}^A$ between segments which are 1^{st} or 2^{nd} -order neighbors. An appearance measure is defined to be:

$$\Delta_{i|j}^A = \|V_i - V_{i \cup j}\|_2.$$

This is the change in the appearance vector of region i caused by merging it with region j . This quantity is asymmetric – i.e. $\Delta_{i|j}^A \neq \Delta_{j|i}^A$. This quantity will encourage merging neighboring regions which have similar appearance vectors.

The appearance similarity measure is modified by an edge-term ($E_{i,j}$ ranging from 0 to 1) that computes the strength of the edge on the boundary between two adjacent regions. This edge term is computed very simply using the Sobel edge detector. The underlying intuition is that we reduce the similarity between adjacent regions if there is an edge between them. We do not introduce an edge-term between segments in the 2^{nd} -order neighborhood (because we want this type of merging to jump between regions) and instead we pay a fixed penalty of size 1 (which is the maximum value the edge term can take).

This gives an asymmetric similarity function $\Delta_{i|j}$:

$$\Delta_{i|j} = \begin{cases} E_{i,j} + \Delta_{i|j}^A & \text{if } i, j \text{ are } 1^{st}\text{-order neighbors} \\ 1 + \Delta_{i|j}^A & \text{if } i, j \text{ are } 2^{nd}\text{-order neighbors} \end{cases}$$

Alternative segmentation algorithms such as the one in [5] can be used. Herein we use the simple method described above in combination with PageRank algorithm or with its quantum extension to rank the pairing between segments (i, j) on the basis of the similarity function. We allow the 30 % highest ranked segments to merge, see Fig. (1). This ranking encourages merging between segments which are most similar. However we reject merges in situations where the similarity function between two regions is too asymmetric (i.e. we do not allow merges where i "likes" j , but j does not "like" i). After these merges, we re-compute the PageRank algorithm and repeat the process.

3 Ranking Based on Random Walks

Since the similarity measure used for merging is asymmetric, we use a directed graph for encoding each segmentation level in the hierarchy.

A directed graph (digraph) $G = (V, E)$ with $N = |V|$ vertices and edges $E \subseteq V \times V$ is encoded by an adjacency matrix \mathbf{A} where $A_{ij} > 0$ if $i \rightarrow j \in E$ and $A_{ij} = 0$ otherwise (this definition includes weighted adjacency matrices). The outdegree matrix D^{out} is a diagonal matrix where $d_i^{out} = \sum_{j \in V} A_{ij}$. The transition matrix \mathbf{P} is defined by $P_{ij} = \frac{A_{ij}}{d_i^{out}}$ if $(i, j) \in E$ and $P_{ij} = 0$ otherwise.

The transition matrix is key to defining random walks on the digraph and P_{ij} is the probability of reaching node j from node i . Given these definitions

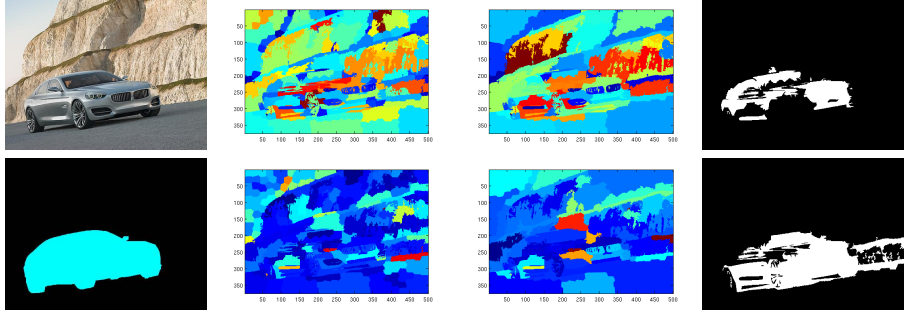


Fig. 2. Quantum vs classical ranking for car segmentation. Top-left: image. Bottom-left: ground truth. Top-center: classical ranking of segments at iterations 8 and 13. Bottom-center: quantum ranking of segments at the same iterations. Top-right: grouping result with the hierarchical algorithm based on the classical ranking (IoU = 0.43). Bottom-right: result using quantum ranking (IoU=0.61).

we have that $\sum_{j \in V} P_{ij} \neq 1$ in general. In addition, \mathbf{P} is irreducible iff G is strongly connected (there is path from each vertex to every other vertex). If P is irreducible, the Perron-Frobenius theorem ensures that there exists a left eigenvector ϕ satisfying $\phi^T \mathbf{P} = \lambda \phi^T$ and $\phi_i > 0 \forall i$. If P is aperiodic (spectral radius $\rho = 1$) we have $\phi^T \mathbf{P} = \rho \phi^T$ and all the other eigenvalues have an absolute value smaller than $\rho = 1$. By ensuring strong connection and aperiodicity we also ensure that any random walk in a directed graph satisfying these two properties converges to a unique stationary distribution.

By correcting \mathbf{P} so that $P_{ij} = \frac{1}{N}$ if $A_{ij} = 0$ and $d_i^{out} = 0$, we obtain a row stochastic matrix: $\sum_{j \in V} P_{ij} = 1 \forall i$. This strategy is adopted in Pagerank [6] and provides and allows for *teleporting* acting on the random walk to any other node in the graph. Teleporting is modeled by defining $\mathbf{G} = \eta \mathbf{P}^T + (1 - \eta) \frac{ee^T}{N}$, where e^T is the all ones row vector and $0 < \eta < 1$. The new matrix G is column stochastic and ensures both irreducibility and aperiodicity. Under these conditions G_{ji} is the probability of reaching j from i . Teleporting means that for every node with $A_{ij} > 0$, $G_{ji} = \frac{A_{ij}}{d_i^{out}}$ is applied with probability η , whereas for nodes with $A_{ij} = 0$ we have $G_{ji} = \frac{1}{N}$ with probability $1 - \eta$. In [7] a trade-off between large values η (preserving more the structure of P') and small ones (potentially increasing the spectral gap) is recommended. For instance, in [3], where the task is to learn classifiers on directed graphs, the setting is $\eta = 0.99$, but usually $\eta = 0.85$ is recommended. In any case, when using the new P we always have that $G_{ii} \neq 0$ due to the Pagerank masking.

Finding the stationary distribution ϕ can be then formulated as an eigenvector problem $\mathbf{G}\phi = \phi$ subject to a normalization constraint $e^T \phi = 1$ (see [8]). Usually the power method is used. Accordingly iterate $\phi(k+1) = \mathbf{G}\phi(k)$ starting by $\phi(0) = e \frac{1}{N}$ until convergence (which will occur if the second eigenvalue λ_2 is smaller than $\lambda_1 = 1$). The stationary distribution is used for ranking the nodes.

4 Quantization of Random Walks

4.1 Unitary vs Stochastic Evolution

The above process for finding the stationary distribution simulates the diffusion of a discrete-time classical random walk on the directed graph $G = (V, E)$. Then the states are the nodes and top-ranked nodes are those whose stationary probability is high.

On the other hand a discrete-time quantum walk [9] diffuses in a very different way since it is subject to quantum superpositions. In this approach states $|\psi\rangle \in \mathbb{C}^N$ are assumed to belong to a Hilbert space $\mathcal{H} = \text{span}\{|j\rangle | j = 1, \dots, N\} = \mathbb{C}^N$ where $\langle j| = (0 \dots 1 \dots 0)$ with a 1 at the j -th position. We use the Dirac bracket notation where: $|a\rangle = \mathbf{a}$, $\langle a| = \mathbf{a}^*$, $\langle a|b\rangle = \mathbf{a}^*\mathbf{b}$ is the inner product and therefore $\langle j|k\rangle = \mathbf{j}^*\mathbf{k} = \delta_{jk}$. Then, the state of the quantum walk at a given time is $|\psi\rangle = \sum_{j=1}^N c_j|j\rangle$ with $c_j \in \mathbb{C}$ so that $|c_1|^2 + |c_2|^2 + \dots + |c_N|^2 = 1$ and $|c_i|^2 = \bar{c}_i c_i$. The probability that the quantum walk is at node i is given by $|\langle i|\psi\rangle|^2 = |c_i|^2$. The $|c_i|^2$ are known as the amplitudes of the wave traveling through the graph.

Given an initial state $|\psi(0)\rangle = \sum_{j=1}^N c_j^0|j\rangle$, a quantum walk evolves through a unitary operator instead of a stochastic one which is the case of random walks do. A $N \times N$ complex matrix \mathbf{U} is unitary if $\mathbf{U}^*\mathbf{U} = \mathbf{U}\mathbf{U}^* = \mathbf{I}_N$, where \mathbf{U}^* is the conjugate transpose, that is $(\mathbf{U}^*)_{ij} = \overline{\mathbf{U}_{ji}}$. Therefore, both the rows and columns of \mathbf{U} form an orthonormal basis in \mathbb{C}^N . In addition \mathbf{U} is by definition a normal matrix for it commutes with its conjugate transpose. In this case it is unitarily similar to a diagonal matrix, i.e., it is diagonalizable by $\mathbf{U} = \mathbf{V}\mathbf{A}\mathbf{V}^*$ where $\mathbf{A} = \text{diag}(\lambda_1 \lambda_2 \dots \lambda_N)$ contains the eigenvalues of \mathbf{U} and \mathbf{V} is unitary and its columns contains the eigenvectors of \mathbf{U} . Combining the latter diagonalization with the property $|\det(\mathbf{U})| = 1$ we have that all the eigenvalues of \mathbf{U} must lie on the unit circle. In other words, they must have either the form $e^{i\theta}$ or the form $e^{-i\theta}$, where θ is an angle on the complex plane.

Therefore we have $|\psi(t)\rangle = \mathbf{U}^t|\psi(0)\rangle$ with the amplitudes of $|\psi(t)\rangle$ summing to unity since \mathbf{U} is unitary.

4.2 Szegedy's Quantization

The problem of associating a unitary operator with a Markov chain (stochastic matrix) has been posed in different ways. One of them is inspired in the Grover's search algorithm [10]. Grover's search relies on projection operators $\mathbf{I} = \sum_{j=1}^N |\Psi_j\rangle\langle\Psi_j|$ where, for instance, $|\Psi_j\rangle = \sum_{k=1}^N \frac{1}{\sqrt{N}}|k\rangle$. The projectors satisfy the condition $\mathbf{I}^2 = \mathbf{I}$ and the operator $2(\mathbf{I} - \mathbf{1})$ defines reflections (coin flips) around the subspace spanned by vectors $|\Psi_j\rangle$. In [11] Szegedy uses a product of reflections for quantizing a Markov chain.

To commence, the state space, originally placed at the N nodes, is moved to the $N \times N$ directed edges of the graph. The Hilbert space is now $\mathcal{H} = \text{span}\{|i\rangle_1|j\rangle_2 : | i, j = 1, \dots, N\} = \mathbb{C}^N \otimes \mathbb{C}^N$ where $|a\rangle|b\rangle = |a, b\rangle = |a\rangle \otimes |b\rangle$ is

the tensor (Kronecker) product, and the subindexes 1 and 2 make explicit the orientation of each edge. Following the approach in [12][13] orientation is critical when computing projections onto the vector encoding the second node of the edge. Then, the superposition of the edges outgoing from node j are given by

$$|\Psi_j\rangle = \sum_{k=1}^N \sqrt{G_{kj}} |j\rangle_1 |k\rangle_2. \quad (1)$$

Given the projector $\Pi = \sum_{j=1}^N |\Psi_j\rangle\langle\Psi_j|$ (coin flip) and the swap operator $\mathbf{S} = \sum_{j=1}^N \sum_{k=1}^N |j\rangle|k\rangle\langle k|\langle j|$ which alternates the direction of the edge (swaps both edge spaces), a step of the quantum walk is given by the unitary operator $\mathbf{U} = \mathbf{S}(2\Pi - \mathbf{1})$, where $\mathbf{1}$ is the identity matrix. However, Grover's search requires a two-step unitary operator per iteration. When translating this idea to Markov chains, Szegedy suggested the use of the operator $(2\Pi - \mathbf{1})(2\Pi' - \mathbf{1})$ in order to contemplate the case of two Markov chains (each one with its own reflection operator). When the two chains are coincident (e.g. for creating bipartite walks) then we have $(2\Pi - \mathbf{1})^2$. If we include the swap operator, which is also unitary, then the two-step evolution operator is given by $\mathbf{U}^2 = (2\mathbf{S}\Pi\mathbf{S} - \mathbf{1})(2\Pi - \mathbf{1})$. This operator swaps the directions of the edges an even number of times

The initial state $|\psi(0)\rangle = \frac{1}{\sqrt{N}} \sum_{j=1}^N \sum_{k=1}^N |j\rangle|k\rangle$ assumes uniform probabilities for all the $N \times N$ edges in the digraph. Given $|\psi(0)\rangle$, we have that the state of the quantum walk at a given time t is given by $|\psi(t)\rangle = \mathbf{U}^{2t}|\psi(0)\rangle$. In addition, the probability of being at such state is $\langle\psi^*(t)|\psi(t)\rangle$. However, for ranking a given node i it is desirable to compute the probability of being at i at time t . This can be done by computing the probability that any edge ends at such node after t time steps. Let the state $|I_q(t)\rangle$ be the superposition of all paths ending at $|i\rangle_2$ (the second space of each edge ending at i).

The superposition state is defined as $|I_i(t)\rangle = {}_2\langle i|\psi(t)\rangle = {}_2\langle i|\mathbf{U}^{2t}|\psi(0)\rangle$ and it is given by the projection of $|\psi(t)\rangle$ onto the space $|i\rangle_2$. Such projection can be described more clearly if we exploit the spectral theorem, since we have $\mathbf{U}^{2t} = \sum_{\mu} \mu^{2t} |\mu\rangle\langle\mu|$, where the μ are the N^2 eigenvalues of \mathbf{U}^2 and the $|\mu\rangle$ are their corresponding N^2 -dimensional eigenvectors. Then, ${}_2\langle i|\mathbf{U}^{2t} = \sum_{\mu} \mu^{2t} {}_2\langle i|\mu\rangle\langle\mu|$. If we consider that the structure of ${}_2\langle i|\mu\rangle$ is ${}_2\langle i|j\rangle_1 |k\rangle_2$, for $|\mu\rangle$ is defined in the tensor space $\mathcal{H} = \mathbb{C}^N \otimes \mathbb{C}^N$, then we have that the proper projection is given by the contraction ${}_2\langle i|k\rangle_2 |j\rangle_1$.

Consequently, we have that the probability that the quantum walk at vertex i after t time steps (that is, its *quantum ranking* at this time) is

$$I_i(t) = \langle\psi(0)|\mathbf{U}^{*2t}|i\rangle_2 \langle i|\mathbf{U}^{2t}|\psi(0)\rangle = \left\| \sum_{\mu} \mu^{2t} {}_2\langle i|\mu\rangle\langle\mu|\psi(0)\rangle \right\|^2 \quad (2)$$

In practice, the time-averaged quantum ranking is used (although the long time average can be used since it only depends on the eigenvectors $|\mu\rangle$). The average is useful because quantum oscillation typically decrease in amplitude with $I_i(t)$. In any case, the main problem when this approach is applied to ranking segments is to compute or to approximate \mathbf{U}^2 (memory storage) and/or the its eigensystem.

4.3 The Eigensystem of the Unitary Operator

One of the nice properties of the Szegedy's formulation is that it provides a direct link between the eigenvalues and eigenvectors of \mathbf{U} and those of a $N \times N$ matrix \mathbf{D} , where $D_{ij} = \sqrt{G_{ij}G_{ji}}$, called the *discriminant matrix*. Such matrix is linked with the projector operator \mathbf{II} and the *half-projection* $\mathbf{A} = \sum_{j=1}^N |\psi_j\rangle\langle j|$ through the following properties: (i) $\mathbf{A}^*\mathbf{A} = \mathbf{1}$, (ii) $\mathbf{A}\mathbf{A}^* = \mathbf{II}$ and (iii) $\mathbf{A}^*\mathbf{S}\mathbf{A} = \mathbf{D}$.

Let λ and $|\lambda\rangle$ be respectively the N eigenvalues and N -dimensional eigenvectors of the symmetric matrix \mathbf{D} . Then, if we define $|\tilde{\lambda}\rangle = \mathbf{A}|\lambda\rangle$ and apply the above properties, we have the following ansatz for the eigenvectors $|\mu\rangle$ and eigenvalues μ of \mathbf{U} : $|\mu\rangle = |\tilde{\lambda}\rangle - \mu\mathbf{S}|\tilde{\lambda}\rangle$. This means that we can easily obtain $2N$ of the N^2 eigenvalues and eigenvectors of \mathbf{U} from the N ones of \mathbf{D} , since $\mu = e^{\pm i \times \arccos \lambda}$. These latter values come from the SVD decomposition of \mathbf{D} , whose singular values lie in $(0, 1)$.

In addition, when we consider \mathbf{U}^2 we have that this operator splits \mathcal{H} into the subspaces $\mathcal{H}_{dyn} = \text{span}\{|\psi_j\rangle, \mathbf{S}|\psi_j\rangle\}$ and its orthogonal complement $\mathcal{H}_{nodyn} = \mathcal{H}_{dyn}^\perp$. The dimension of \mathcal{H}_{dyn} is at most $2N$. Thus, the spectrum of \mathbf{U}^2 corresponding to \mathcal{H}_{dyn} is given by, at most, the $2N$ values $\{e^{\pm 2i \times \arccos \lambda}\}$. The spectrum corresponding to \mathcal{H}_{nodyn} is given by at least $N^2 - 2N$ 1's.

When estimating the eigenvectors $|\mu\rangle$ and eigenvalues μ for segmentations we have to confront the problem that at the lowest segmentation levels the number of SLIC superpixels is too high for building an $N^2 \times N^2$ unitary operator and then getting its eigensystem. The unitary operator is needed to extract the eigenvectors $|\mu\rangle$ corresponding to eigenvectors with value 1 (i.e. satisfying $\mathbf{U}^2|\mu\rangle = |\mu\rangle$). Obtaining the operator \mathbf{U}^2 (or even \mathbf{U}) is infeasible for these levels, unless a more in-depth analysis of the structure of these operators reveals a shortcut. This latter question is beyond the scope of this paper and we have approximated the instantaneous ranking $I_i(t)$ with

$$\tilde{I}_i(t) = \left\| \sum_{\mu \in \mathcal{H}_{dyn}} \mu^{2t} {}_2\langle i|\mu\rangle\langle\mu|\psi(0)\rangle \right\|^2. \quad (3)$$

Therefore $\tilde{I}_i(t)$ is a low-pass approximation of $I_i(t)$.

5 Analysis: Experiments and Conclusions

We evaluate to what extent the averaged $\tilde{I}_i(t)$ can improve grouping when applied to rank segments at all levels of the hierarchy. In order to do that, we measure the Intersection-over-Union (IoU) which quantifies this quality of the segmentation of a particular object class (Pascal VOC 2010). This measure penalizes both obtaining a smaller area than the ground truth and obtaining a larger area than the ground truth. We have access to an unpublished ground truth where it has been assigned one of 57 labels to every pixel. For quantification data we use a sample of 1,100 images of the 10,103 in the VOC 2010

	plane	bicycle	bird	boat	bottle	bus	car	cat
CPMC	78.6	64.4	74.8	71.2	74.9	78.9	72.1	85.5
Uniform	53.7	39.7	55.8	48.3	54.6	46.5	51.7	57.1
R. Walks	53.3	41.9	57.8	56.9	51.7	52.5	53.9	60.5
Quantum	55.6	41.8	56.3	54.7	51.7	49.0	53.0	58.0
	chair	cow	d. table	dog	horse	motorbike	person	pot. plant
CPMC	51.7	79.9	62.2	82.6	77.5	74.2	67.6	63.9
Uniform	52.1	54.9	57.1	55.9	54.1	52.5	48.4	50.1
R. Walks	52.9	57.0	48.6	57.3	54.6	48.9	51.1	51.7
Quantum	49.7	57.1	51.0	57.5	51.8	49.2	50.3	51.6
	sheep	sofa	train	tv	bag	bed	bench	book
CPMC	73.7	64.2	78.2	76.8	53.2	65.4	36.5	36.3
Uniform	54.6	58.2	48.0	57.9	54.3	54.3	43.1	46.8
R. Walks	56.8	58.8	50.9	57.2	55.7	59.4	49.0	52.9
Quantum	56.1	57.5	48.5	57.5	54.9	62.4	54.6	55.1
	building	cabinet	ceiling	clothes	pc	cup	door	fence
CPMC	51.6	54.9	22.4	58.5	70.7	40.9	40.0	42.3
Uniform	51.8	55.5	54.6	57.3	48.8	47.3	50.8	44.5
R. Walks	53.5	54.6	60.1	57.8	48.6	46.8	56.2	45.7
Quantum	55.1	59.0	51.0	58.8	52.0	51.2	59.4	47.4
	floor	flower	food	grass	ground	keyboard	light	mountain
CPMC	54.5	51.0	44.9	56.5	55.0	46.8	8.3	55.2
Uniform	58.8	49.8	43.3	60.4	58.5	46.3	43.7	59.5
R. Walks	62.1	48.5	56.9	61.3	60.4	52.2	38.4	62.3
Quantum	65.0	49.9	50.7	61.8	60.9	53.5	42.5	63.8
	mouse	sign	plate	road	rock	shelves	sidewalk	sky
CPMC	12.6	24.6	44.5	56.3	61.1	52.6	53.0	65.1
Uniform	48.2	39.1	43.6	64.8	58.2	47.8	57.2	76.4
R. Walks	48.4	44.9	39.8	66.5	55.3	54.0	61.9	77.2
Quantum	48.3	47.4	39.0	67.5	58.2	49.9	59.2	81.0
	snow	table	track	tree	truck	wall	water	window
CPMC	60.0	48.9	42.3	54.5	61.8	51.7	65.4	50.9
Uniform	55.8	47.0	41.8	54.2	47.1	59.0	62.4	54.2
R. Walks	55.1	49.2	44.0	56.3	53.7	59.7	62.9	57.9
Quantum	59.1	47.2	41.7	56.7	55.8	61.1	65.1	56.1
	wood	all IoU						
CPMC	49.4	59.6						
Uniform	59.2	57.2						
R. Walks	60.5	58.7						
Quantum	60.3	59.4						

Table 1. IoU wrt the ground truth for 57 object classes of the Pascal VOC 2010

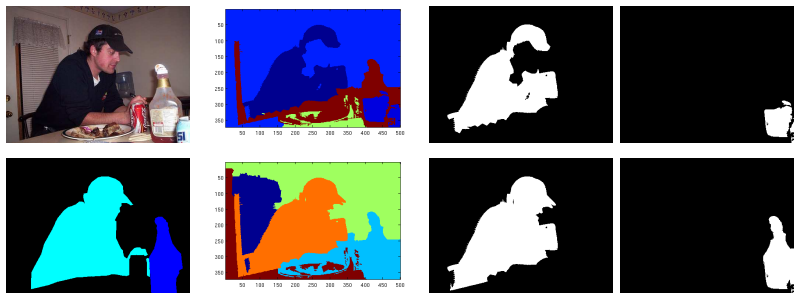


Fig. 3. Quantum vs classical ranking for person and bottle segmentation. First row, from left to right: image, classical rank at iteration 28, grouping result using classical rank for person (IoU=0.65), grouping using classical rank for bottle (IoU=0.36) . Second row, from left to right: ground truth, quantum rank at iteration 28, result using quantum rank for person (IoU=0.80), result with quantum rank for bottle (IoU=0.67).

dataset. In all cases the number of initial SLIC super pixels is 500 and there are 35 levels in the hierarchy until we reach a single segment.

Firstly, we analyze the behavior of Pagerank vs qrank (its low-pass approximation), before going through IoU analysis. In Fig. 2 we try to detect a car embedded in a textured environment. Pagerank seems to invert the priorities of qrank. It prioritizes the selection of segments inside quasi-homogeneous regions, whereas boundary segments have a low rank (blue). However, once the homogeneous region is built, its ranking decreases (see the grouping of the sky, in first row, which starts at iteration 8 and it is stopped at iteration 13). Most top-ranked segments in qrank correspond to low-ranked ones in Pagerank (see the shadowed region at the top of the car). However, this "inversion" is misleading since some top/medium-ranked segments in qrank are also prioritized by Pagerank. On the one hand, the tendency of Pagerank to propose merging inside a quasi-homogeneous region may merge a part of the object with the background when the appearance of the background and the part of the object are similar. In the case of the car, this results in obtaining a smaller area than the ground truth. On the other hand, the behavior of qrank may lead us to merge the object with a part of the background producing a bigger area than the ground truth. These behaviors are replicated in Fig. 3 where qrank outperforms Pagerank when segmenting a person and a bottle.

We compare the IoUs obtained by both ranking methods with: (i) a state-of-the-art unsupervised segmentation algorithm (CPMC) and (ii) our hierarchical method with uniform ranking probabilities. We show the quantitative results in Table. 1. In general CPMC [14] outperforms our method. The first stage of the Constrained Parametric Min-Cuts method (CPMC) applies repeatedly a max-flow algorithm to output a set of segments based on groupings of an edge map provided by the gPb algorithm [15]. After a non-maximum suppression filtering, the second stage ranks the segments using cues trained on the Pascal VOC dataset, provided the groundtruth masks of the objects.

However, our grouping method uses only low-level cues and is not trained for object-like segments. Thus, it is natural to outperform CPMC on background region classes, for which it is not trained. Also, our method tends to group similar appearance regions, while CPMC may group different appearance regions. This explains that CPMC outperforms our method on many foreground object classes and on the overall performance.

To conclude, qrank slightly outperforms its classical counterpart, and in some cases both outperform a state-of-the-art learning-based method. These experiments provide an initial insight of the power of quantum walks but we must complete the low-pass ranking with higher-order experiments which become feasible in segmentation settings.

Acknowledgements Funding. F. Escolano: Project TIN2012-32839 (Spanish Gov.). E. R. Hancock: Royal Society Wolfson Research Merit Award.

References

1. Grady, L.: Random walks for image segmentation. *TPAMI* **28**(11) (2006) 1768–1783
2. Burges, C.J.C., Platt, J.C.: Semi-supervised learning with conditional harmonic mixing. In Chapelle, O., Schölkopf, B., Zien, A., eds.: *Semi-Supervised Learning*, Cambridge, MA, MIT Press (2006)
3. Zhou, D., Huang, J., Schölkopf, B.: Learning from labeled and unlabeled data on a directed graph. In: *ICML*. (2005) 1041–1048
4. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. *TPAMI* **34**(11) (2012) 2274–2282
5. Nock, R., Nielsen, F.: Statistical region merging. *TPAMI* **26**(11) (2004) 1452–1458
6. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web (1999)
7. Johns, J., Mahadevan, S.: Constructing basis functions from directed graphs for value function approximation. In: *ICML*. (2007) 385–392
8. Langville, A.N., Meyer, C.D.: Deeper inside pagerank. *Internet Mathematics* **1** (2004) 335–400
9. Aharonov, Y., Davidovich, L., Zagury, N.: Quantum random walks. *Phys. Rev. A* **48** (August 1993) 1687–1690
10. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: *ACM Symposium on Theory of Computing*. (1996) 212–219
11. Szegedy, M.: Quantum speed-up of markov chain based algorithms. In: *FOCS*, IEEE Computer Society (2004) 32–41
12. Paparo, G.D., Martin-Delgado, M.A.: Google in a quantum network. *CoRR abs/1112.2079* (2011)
13. Paparo, G.D., Müller, M., Comellas, F., Martin-Delgado, M.A.: Quantum google in a complex network. *Scientific Reports* **3:2773** (2013)
14. Carreira, J., Sminchisescu, C.: CPMC: Automatic object segmentation using constrained parametric min-cuts. *TPAMI* **34**(7) (2012) 1312–1328
15. Maire, M., Arbelaez, P., Fowlkes, C.C., Malik, J.: Using contours to detect and localize junctions in natural images. In: *CVPR*. (2008)