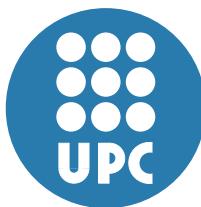Developing Competitive HMM PoS Taggers Using Small
Training Corpora
Padró, M. and Padró, Ll.
Research Report LSI-04-36-R

Developing Competitive HMM PoS Taggers Using Small Training Corpora

Muntsa Padró and Lluís Padró
TALP Research Center
Universitat Politècnica de Catalunya
Jordi Girona Salgado 1-3
08034 Barcelona
{mpadro,padro}@lsi.upc.es

February 2004

Abstract

This paper presents a study aiming to find out the best strategy to develop a fast and accurate HMM tagger when only a limited amount of training material is available. This is a crucial factor when dealing with languages for which small annotated material is not easily available.

First, we develop some experiments in English, using WSJ corpus as a test-bench to establish the differences caused by the use of large or a small train set. Then, we port the results to develop an accurate Spanish PoS tagger using a limited amount of training data.

Different configurations of a HMM tagger are studied. Namely, trigram and 4-gram models are tested, as well as different smoothing techniques. The performance of each configuration depending on the size of the training corpus is tested in order to determine the most appropriate setting to develop HMM PoS taggers for languages with reduced amount of corpus available.

1   Introduction

PoS Tagging is a need for most of Natural Language applications such as Sumarization, Machine Translation, Dialogue systems, etc. and the basis of many higher level NLP processing tasks. It is also used to obtain annotated corpora combining automatic tagging with human supervision. These corpora may be used for linguistic research, to build better taggers, or as statistical evidence for many other language-processing related goals.

PoS tagging has been largely studied and many systems developed. There are some statistical implementations such as those presented in [1, 2, 3, 4] and some knowledge-based taggers (finite-state, rule-based, memory based) [5, 6, 7]. There are also some systems that combine different implementations with a voting procedure.

This work presents a thorough study aiming to establish which is the most appropiate way to train a HMM PoS tagger when dealing with languages with a limited amount of training corpora. To do so, we compare different smoothing techniques and different order HMMs.

Experiments are performed to determine the performance of the best configuration when the tagger is trained with a large English corpus (1 million words from WSJ), and comparing the results with those for a training corpus ten times smaller. Then, the experiments for the small train corpus are repeated in another language (Spanish), validating the conclusions. The tested HMM configurations vary on the order of the model (3 and 4 order HMMs are tested) and in the smoothing technique (Lidstone's law *vs.* Linear Interpolation) used to estimate model parameters.

Section 2 presents the theoretical basis of a HMM and the different smoothing techniques used in this work. Section 3 shows the realized experiments and the obtained results. Section 4 states some conclusions and further work.

2   Hidden Markov Models

We will be using Hidden Markov Models Part-of-Speech taggers of order three and four. Depending on the order of the model, the states represent pairs or triples of tags, and obviously, the number of parameters to estimate varies largely. The parameters of such models are initial state probabilities, state transition probabilities and emission probabilities. That is:

$$\pi_i = P(q_1 = s_i)$$

is the probability that a sequence starts at state s_i,

$$a_{ij} = P(q_{t+1} = s_j | q_t = s_i)$$

is the transition probability from the state i to the state j (i.e. trigram probability $P(t_3|t_1 t_2)$ in a $3^r d$ order model, or 4-gram probability $P(t_4|t_1 t_2 t_3)$ in a four-gram HMM), and

$$b_i(k) = P(w_k | s_i)$$

is the emission probability of the symbol w_k from state s_i.

In the PoS task model, the emitted symbols are the observed words and the observed sequence is a complete sentence. Given a sentence, we want to choose the most likely sequence of states (i.e. PoS tags) that generated it. This is computed using the Viterbi algorithm [8].

2.1 Parameter Estimation

The simplest way to estimate the HMM parameters is Maximum Likelihood Estimation (MLE), which consists in computing the relative frequency of each observation:

$$P_{MLE}(x) = \frac{count(x)}{N}$$

$$P_{MLE}(x|y) = \frac{count(x,y)}{count(x)}$$

For the case of the HMM we have to compute this probability estimation for each initial state, transition or emission in the training data:

$$\pi_i = \frac{count(s_i(t=0))}{count(sentences)}$$

$$a_{ij} = \frac{count(s_i \to s_j)}{count(s_i)}$$

$$b_i(k) = \frac{count(s_i, w_k)}{count(s_i)}$$

Actually, computing $b_i(k)$ in this way is quite difficult because the number of occurrences of a single word will be too small to provide enough statistical evidence, so Bayes rule is used to compute $b_i(k)$ as:

$$b_i(k) = P(w_k|s_i) = \frac{P(s_i|w_k)P(w_k)}{P(s_i)}$$

where:

$$P(s_i) = \frac{count(s_i)}{count(words)}$$

$$P(w_k) = \frac{count(w_k)}{count(words)}$$

Since $P(s_i|w_k)$ would also require lots of data to be properly estimated, we approximate it as $P(t|w_k)$, where t is the last tag in the n-gram corresponding to the state. Similarly, $P(s_i)$ is approximated as $P(t)$.

2.2 Smoothing

MLE is usually a bad estimator for NLP purposes, since data tends to be sparse[1]. This leads to zero probabilities being assigned to unseen events, causing troubles when multiplying probabilities.

To solve this sparseness problem it is necessary to look for estimators that assign a part of the probability mass to the unseen events. To do so, there are many different smoothing techniques, all of them consisting of decreasing the probability assigned to the seen events and distributing the remaining mass among the unseen events. In this work two smoothing methods are compared: Lidstone's law and Linear Interpolation.

2.2.1 Laplace and Lidstone's Law

The oldest smoothing technique is Laplace's law [9], that consists in adding one to all the observations. That means that all the unseen events will have their probability computed as if they had appeared once in the training data. Since one observation for each event (seen or unseen) is added, the number of different possible observations (B) has to be added to the number of real observations (N), in order to maintain the probability normalised.

$$P_{La}(x) = \frac{count(x + 1)}{N + B}$$

However, if the space is large and very sparse –and thus the number of possible events (B) may be similar to (or even larger than) the number of observed events– Laplace's law gives them too much probability mass.

A possible alternative is Lidstone's law (see [10] for a detailed explanation on these and other smoothing techniques) which generalises Laplace's and allows to add an arbitrary value to unseen events. So, for a relatively large number of unseen events, we can choose to add values lower than 1. For a relatively small number of unseen events, we may choose 1, or even larger values, if we have a large number of observations (N).

$$P_{Ld}(x) = \frac{count(x) + \lambda}{N + B\lambda} \quad \lambda > 0$$

To use Laplace's or Lidstone's laws in a HMM-based tagger we have to smooth all probabilities involved in the model:

$$\pi_i = \frac{count(s_i(t = 0)) + \lambda_\pi}{count(sentences) + B_{tag}\lambda_\pi}$$

---

[1]Following Zipf's law: a word's frequency is inversely proportional to its rank order

$$a_{ij} = \frac{count(s_i \to s_j) + \lambda_A}{count(s_i) + B_{tag}\lambda_A}$$

$$P(s_i) = \frac{count(s_i) + \lambda_s}{count(words) + B_{tag}\lambda_s}$$

$$P(w_k) = \frac{count(w_k) + \lambda_w}{count(words) + B_w\lambda_w}$$

where B_{tag} is the number of possible tags and B_w is the number of words in the vocabulary (obviously, we can only approximate this quantity).

Since there are different counts involved in each probability, we have to consider different λ values for each formula. In the case of Laplace's law, all λ are set to 1, but when using Lidstone's, we want to determine which is the best set of λ, and how they vary depending on the train set size. This is further discussed in section 3.1.

2.2.2 Linear Interpolation

A more sophisticated smoothing technique consists of linearly combine the estimations for different order n-grams:

$$
\begin{aligned}
P_{li}(t_3|t_1t_2) &= c_1P(t_3) + c_2P(t_3|t_2) + c_3P(t_3|t_1t_2) \\
P_{li}(t_4|t_1t_2t_3) &= c_1P(t_4) + c_2P(t_4|t_3) + c_3P(t_4|t_2t_3) + c_4P(t_4|t_1t_2t_3)
\end{aligned}
$$

where $\sum_i c_i = 1$ to normalise the probability. Although the values for c_i can be determined in many different ways, in this work they are estimated by deleted interpolation as described in [4]. This technique assumes that the c_i values don't depend on the particular n-gram and computes the weights depending on the counts for each i-gram involved in the interpolation.

3 Experiments and Results

The main purpose of this work is to study the behaviour of different configurations for a HMM-based PoS tagger, in order to determine the best choice to develop taggers for languages with scarce annotated corpora available.

First, we will explore different configurations when a large amount of training corpus is available. The experiments will be performed on English, using 1.1 million words from the Wall Street Journal corpus. Then, the same configurations will be explored when the training set is reduced to 100,000 words.

Later, the behaviour on the reduced train set will be validated on a manually developed 100,000 word corpus for Spanish [11].

The tested configurations vary on the order of the used HMM (trigram or 4-gram), and on the smoothing technique applied (Lidstone's law or Linear Interpolation).

All the experiments are done using ten fold cross-validation. In each fold, 90% of the corpus is used to train the tagger and the rest to test it.

In the following sections, we present the behaviour of the different HMM configurations for English, both with a large and a small corpus. After, we repeat the tests using a small corpus for Spanish.

3.1   Applying Lidstone's Law

As it has been explained in section 2.2.1, when Lidstone's law is used in a HMM tagger, there are four λ values to consider. Changing these values significantly affects the precision of the system.

Thus, before comparing this smoothing technique with another, we have to select the set of λ that yields the best tagger performance. After performing some experiments, we observed that λ_A is the only parameter that significantly affects the behaviour of the system. Modifying the other three values didn't change the system precision in a significant way. So λ_π, λ_s, and λ_w were set to some values determined as follows:

- λ_π is the assigned count for unobserved initial states. Since initial states depend only on the tag of the first word in the sentence, and the tag set we are using is quite reduced (about 40 tags), we may consider that in a 1,000,000 word corpus, at least one sentence will start with each tag. So, we will count one occurrence for unseen events (i.e. we are using $\lambda_\pi = 1$, Laplace's law, for this case). When the corpus is ten times smaller, we will use a proportional rate of occurrence ($\lambda_\pi = 0.1$).

- λ_s is the count assigned to the unseen states. Since we approximate $P(s_i)$ by $P(t)$ (see section 2.1), the possible events are the number of tags in the tag set, and we can reason as above, assuming at least one occurrence of each tag in a 1,000,000 word corpus (again, Laplace's law, $\lambda_s = 1$), and a proportional value for the small corpus ($\lambda_s = 0.1$)

- λ_w is the count assigned to the unseen words. Obviously, enforcing that each possible word will appear at least once would take too many probability mass out of seen events (English vocabulary is about 100,000 forms, which would represent 10% of a 1 million word corpus), so we adopt a more conservative value: $\lambda_w = 0.1$ for the large corpus, and the proportional value $\lambda_w = 0.01$ for the small one.

After setting these three λ values, we have to select the best value for λ_A. We will repeatedly use hill-climbing with different starting values and step lengths ($\Delta\lambda$), and choose the value that produces better results.

In table 1 the results of this hill-climbing algorithm using the whole English corpus (1 million of words) are presented. Table 2 show the same results for the 100.000 words English corpus.

Initial λ_A	$\Delta\lambda$	Trigram HMM			4-gram HMM		
		Initial precision	Final		Initial precision	Final	
			precision	λ_A		precision	λ_A
0.05	0.01	96.98	97.00	0.22	96.72	96.88	0.28
0.05	0.005	96.98	96.99	0.085	96.72	96.81	0.125
0.5	0.1	97.008	97.009	0.6	96.91	96.93	1.0
0.5	0.05	97.008	97.009	0.4	96.91	96.93	0.95
1.0	0.5	97.00	**97.01**	0.5	96.93	**96.94**	1.5
1.0	0.1	97.00	**97.01**	0.6	96.93	96.93	1.0

Table 1: Precision obtained with the hill-climbing algorithm on the complete English corpus

Initial λ_A	$\Delta\lambda$	Trigram HMM			4-gram HMM		
		Initial precision	Final		Initial precision	Final	
			precision	λ_A		precision	λ_A
0.05	0.01	96.56	96.63	0.09	95.79	96.30	0.33
0.05	0.005	96.56	96.63	0.1	95.79	96.20	0.2
0.5	0.1	96.69	96.69	0.5	96.36	96.43	0.8
0.5	0.05	96.69	96.69	0.5	96.36	96.43	0.75
1.0	0.5	96.70	96.70	1.0	96.44	**96.51**	3.5
1.0	0.1	96.70	**96.71**	0.9	96.44	96.46	1.2

Table 2: Precision obtained with the hill-climbing algorithm on the reduced English corpus

As it was expected, when using a small corpus the precision falls, specially when a 4-gram HMM is used, since the evidence to estimate the model is insufficient. This point is discussed in section 3.3.

These results show that the value selected for λ_A is an important factor when using this smoothing technique. As can be seen in table 2, the precision of the tagger varies up 0.7% depending on the value used for λ_A.

After performing the hill climbing search, we choose the λ_A that gives better results in each case, as the optimal parameter to use with this smoothing technique. So, for the whole corpus using a Trigram HMM, λ_A is set to 0.6 and the tagger yields a precision of 97.01%, while if we use a 4-gram HMM, $\lambda_A = 1.5$ leads to a precision of 96.94%. When the experiments are performed over the reduced corpus, the best results are obtained with $\lambda_A = 0.9$ for a trigram HMM (96.71%) and with $\lambda_A = 3.5$ for a 4-gram model (96.51%).

3.2 Applying Linear Interpolation

The performance of the taggers when using Linear Interpolation to smooth the probability estimations has been also tested. In this case, the coefficients c_i are found via the deleted interpolation algorithm (see section 2.2.2).

When using Linear Interpolation, the precision obtained by the system with the whole corpus is 97.00% with a trigram HMM, and 97.02% with a 4-gram HMM. For the reduced corpus the precision falls slightly and we obtain 96.84% for the trigram model and 96.71% for the 4-gram HMM.

The results obtained using Linear Interpolation and a trigram model should reproduce those reported by [4], where the maximum precision reached by the system on WSJ is 96.7%. In our case we obtain a higher precision because we are assuming the nonexistence of unknown words (i.e. the dictionary contains all possible tags for all words appearing in the test set. Obviously, word-tag frequency information from the test corpus is not used when computing $P(s_i|w_k)$).

3.3 Best Configuration for English

Best results obtained for each HMM tagger configuration are summarized in table 3. Results are given both for the large and small corpus.

1,1 Mword English corpus

	Lidstone's law	Linear Interpolation
trigram	97.01%	97.00%
4-gram	96.94%	**97.02%**

100 Kword English corpus

	Lidstone's law	Linear Interpolation
trigram	96.71%	**96.84%**
4-gram	95.51%	96.71%

Table 3: Obtained results for all HMM PoS tagger configurations using large and small sections of WSJ corpus

Comparing the results for the two smoothing methods used with different order models, we can draw the following conclusions:

– In general, Linear Interpolation produces taggers with higher precision than using Lidstone's law.

– For the case of the large corpus, the results are not significantly different for any combination of n-gram order and smoothing technique. While for the reduced corpus it is clearly better to use a trigram model than a 4-gram HMM, and Linear Interpolation yields slightly better results.

– Using Linear Interpolation has the benefit that the involved coefficients are computed using the training data via deleted interpolation, while for Lidstone's law the precision is very dependent on the λ_A value, which has to be costly optimised (e.g. via hill-climbing).

3.4 Behaviour in Spanish

The same experiments performed for the English corpus were performed with a Spanish corpus (the CLiC-TALP Corpus[2]) which has about 100,000 words. This corpus is manually validated so, although it is small, it is more accurately tagged than WSJ.

In this case the tagger relies on FreeLing morphological analyser [12] instead of using a dictionary built from the corpus. Nevertheless, the situation is comparable to the English experiments above: Since the corpus and the morphological analyser have been hand-developed and cross-checked, they are mutually consistent, and so we don't have to care about unknown words in the test corpus.

3.4.1 Applying Lidstone's Law

In the same way than for the English corpus, a hill-climbing search is performed to study the influence of λ_A value in the precision of the system. The λ_π, λ_s and λ_w values are fixed to the same values used for the reduced WSJ.

Table 4 shows the results of these experiments. The best λ_A for the trigram HMM is $\lambda_A = 1.5$, yielding a precision of 96.85%. The best value for a 4-gram model is $\lambda_A = 2.5$, which produces a precision of 96.22%

---

[2]more information in `http://www.lsi.upc.es/~nlp/`

		Trigram HMM			4-gram HMM		
Initial λ_A	$\Delta\lambda$	Initial precision	Final precision	λ_A	Initial precision	Final precision	λ_A
0.05	0.01	96.54	96.68	0.18	95.49	96.00	0.35
0.05	0.005	96.54	96.58	0.065	95.49	95.82	0.15
0.5	0.1	96.79	96.80	0.5	96.06	96.16	1.6
0.5	0.05	96.79	96.81	0.55	96.06	96.14	1.05
1.0	0.5	96.80	**96.85**	1.5	96.14	**96.22**	2.5
1.0	0.1	96.80	96.84	1.1	96.14	96.16	1.6

Table 4: Precision obtained with the hill-climbing algorithm on the Spanish corpus

3.4.2   Applying Linear Interpolation

The coefficients for Linear Interpolation are computed for Spanish in the same way than for English (section 3.2). The precision of the obtained taggers is 96.90% for the trigram HMM and 96.73% for the 4-gram model.

3.4.3   Best Configuration for Spanish

Results for Spanish are –as it may be expected– similar to those obtained with the reduced English corpus. Again, working with a trigram HMM gives higher precision than working with a 4-gram one, for both smoothing techniques, and using Linear Interpolation gives a slightly better results than using Lidstone's law. Table 5 summarizes the obtained results for both smoothing methods.

100 Kword Spanish corpus

	Lidstone's law	Linear Interpolation
trigram	96.85%	**96.90%**
4-gram	96.22%	96.73%

Table 5: Obtained results for all HMM PoS tagger configurations using Spanish 100 Kwords corpus

Nevertheless, some important remarks can be extracted from these results:

– Competitive HMM taggers may be build using relatively small train sets, which is interesting for languages lacking large resources.

– The best results are obtained using trigram models and Linear Interpolation smoothing.

– Lidstone's law may be used as an alternative smoothing technique, but if λ_A is not tuned, results are likely to be significantly lower.

4   Conclusions and Further Work

We have studied how competitive HMM-based PoS taggers can be developed using relatively small training corpus.

Results point out that accurate taggers can be build provided the appropriate smoothing techniques are used. Between both techniques studied here, in general the one that gives a higher precision is Linear Interpolation but Lidstone's law can reach, in many cases, similar precision rates if a search is performed through the parameter space to find the most appropriate λ_A.

The model proposed in [4] (trigram tagger, Linear Interpolation smoothing) is not only the more suitable for big training corpus but also it gives the best results for limited amounts of training data.

The use of four-gram models may result in a slight increase in precision when using large corpus. Nevertheless, the gain is probably not worth the increase in complexity and size of the model.

Further work to be performed includes dealing with unknown words, and study their influence on the taggers developed on small corpus. Also, we plan to port the same experiments to other languages (namely, Catalan) to further validate the conclusions of this paper.

References

[1] Cutting, D., Kupiec, J., Pedersen, J.O., Sibun, P.: A practical part–of–speech tagger. In: Proceedings of the 3rd Conference on Applied Natural Language Processing, ANLP, ACL (1992) 133–140

[2] Schmid, H.: Improvements in part–of–speech tagging with an application to german. In: Proceedings of the EACL SIGDAT Workshop, Dublin, Ireland (1995)

[3] Ratnaparkhi, A.: A maximum entropy part–of–speech tagger. In: Proceedings of the 1st Conference on Empirical Methods in Natural Language Processing, EMNLP. (1996)

[4] Brants, T.: Tnt - a statistical part- of-speech tagger. In: Proceedings of the 6th Conference on Applied Natural Language Processing, ANLP, ACL (2000)

[5] Karlsson, F.: Constraint grammar as a framework for parsing running text. In: Proceedings of 13th International Conference on Computational Linguistics, COLING, Helsinki, Finland (1990) 168–173

[6] Brill, E.: A Corpus–based Approach to Language Learning. PhD thesis, Department of Computer and Information Science, University of Pennsylvania (1993) http://www.cs.jhu.edu/~brill/acadpubs.html.

[7] Daelemans, W., Zavrel, J., Berck, P., Gillis, S.: Mbt: A memory–based part–of–speech tagger generator. In: Proceedings of the 4th Workshop on Very Large Corpora, Copenhagen, Denmark (1996) 14–27

[8] Viterbi, A.J.: Error bounds for convolutional codes and an asymptotically optimal decoding algorithm. IEEE Transactions on Information Theory (1967) 260–269

[9] Laplace, P.S.m.: Philosophical Essay on Probabilities. Springer-Verlag (1995)

[10] Manning, C.D., Schütze, H.: Foundations of Statistical Natural Language Processing. The MIT Press (1998)

[11] Civit, M.: Criterios de etiquetación y desambiguación morfosintáctica de corpus en español. PhD thesis, Linguistics Department, Universitat de Barcelona (2003)

[12] Carreras, X., Chao, I., Padró, L., Padró, M.: Freeling: An open-source suite of language analyzers. In: Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC'04), Lisbon, Portugal (2004)

Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya

Research Reports - 2004

- LSI-04-1-R : *Automatic Generation of Polynomial Loop Invariants: Algebraic Foundations*, Rodríguez, E. and Kapur, D.

- LSI-04-2-R : *Comparison of Methods to Predict Ozone Concentration* , Orozco, J.

- LSI-04-3-R : *Towards the definition of a taxonomy for the cots product´s market* , Ayala, Claudia P.

- LSI-04-4-R : *Modelling Coalition Formation over Time for Iterative Coalition Games*, Mérida-Campos, C. and Willmott, S.

- LSI-04-5-R : *Illegal Agents? Creating Wholly Independent Autonomous Entities in Online Worlds*, Willmott, S.

- LSI-04-6-R : *An Analysis Pattern for Electronic Marketplaces*, Queralt, A. and Teniente, E.

- LSI-04-7-R : *Exploring Dopamine-Mediated Reward Processing through the Analysis of EEG-Measured Gamma-Band Brain Oscillations*, Vellido, A. and El-Deredy, W.

- LSI-04-8-R : *Studying Embedded Human EEG Dynamics Using Generative Topographic Mapping*, Vellido, A. and El-Deredy, W. and Lisboa, P.J.G.

- LSI-04-9-R : *Similarity and Dissimilarity Concepts in Machine Learning*, Orozco, J.

- LSI-04-10-R : *A Framework for the Definition of Metrics for Actor-Dependency Models*, Quer, C. and Grau, G. and Franch, X.

- LSI-04-11-R : *QM: A Tool for Building Software Quality Models*, Carvallo, J.P. and Franch, X. and Grau, G. and Quer, C.

- LSI-04-12-R : *COSTUME: A Method for Building Quality Models for Composite COTS-based Software Systems*, Carvallo, J.P. and Franch, X. and Grau, G. and Quer, C.

- LSI-04-13-R : *Enabling Collaboration in Virtual Reality Navigators*, Theoktisto, V. and Fairén, M. and Navazo, I.

- LSI-04-14-R : *DesCOTS: A Software System for Selecting COTS Components*, Carvallo, J.P. and Franch, X. and Grau, G. and Quer, C.

- LSI-04-15-R : *Evaluation and symmetrisation of alignments obtained with the Giza++ software*, Lambert, P. and Castell, N.

- LSI-04-16-R : *A note on the use of topology extensions for provoking instability in communication networks*, Blesa, M.J.

- LSI-04-17-R : *An ISO/IEC-compliant Quality Model for ER Diagrams*, Costal, D. and Franch, X.

- LSI-04-18-R : *A Case Study on Pruning General Ontologies for the Development of Conceptual Schemas* , Conesa, J.

- LSI-04-19-R : *Adding Efficient and Reliable Access Paths to the JCF*, Marco, J. and Franch, X.

- LSI-04-20-R : *Exploiting Simple Corporate Memory in Iterative Coalition Games*, Mérida-Campos, C. and Willmott, S.

- LSI-04-21-R : *On the Semantics of Operation Contracts in Conceptual Modeling* , Queralt, A. and Teniente, E.

- LSI-04-22-R : *Complexity issues on bounded restrictive H-coloring*, Díaz, J. and Serna, M. and Thilikos, D.M.

- LSI-04-23-R : *Chromatic number in random scaled sector graphs*, Díaz, J. and Sanwalani, V. and Serna, M. and Spirakis, P.

- LSI-04-24-R : *Bounds on the bisection width for random d-regular graphs*, Díaz, J. and Serna, M. and Wormald, N.C.

- LSI-04-25-R : *Open Source environment to define constraints in route planning for GIS-T*, Pérez, L. and Silveira, A. da M.

- LSI-04-26-R : *A basic repository of operations for the refinement of general ontologies*, de Palol, X.

- LSI-04-27-R : *Tetrahedral mesh subdivision based on underlying volume data*, Rodríguez, L. and Navazo, I. and Vinacua, A.

- LSI-04-28-R : *The Price of Connectedness in Expansions*, Fomin, F.V. and Fraigniaud, P. and Thilikos, D.M.

- LSI-04-29-R : *Smaller kernels for hitting set problems of constant arity*, Nishimura, N. and Ragde, P. and Thilikos, D.M.

- LSI-04-30-R : *Searching Spatial Sense in the Ontological World: Discovering Spatial Objects*, Morocho, V and Pérez, L. and Saltor, F.

- LSI-04-32-R : *Implementation considerations of an Expert System to assess Stream Water Quality management*, Cabanillas, D. and Willmott, S.

- LSI-04-33-R : *Multisided patches*, Pla, N. and Vigo, M. and Cotrina, J.

- LSI-04-34-R : *SVMTool: A general POS tagger generator based on Support Vector Machines*, Giménez, J. and Màrquez, Ll.

- LSI-04-35-R : *A distributed and mobile component system based on the ambient calculus*, Mylonakis, N. and Orejas, F.

- LSI-04-36-R : *Developing Competitive HMM PoS Taggers Using Small Training Corpora*, Padró, M. and Padró, Ll.

Hardcopies of reports can be ordered from:

Campus Nord, Mòdul C6
Jordi Girona Salgado, 1-3
03034 Barcelona,Spain
nurias@lsi.upc.es

See also the Departament WWW pages, http://www.lsi.upc.es/