On Deriving Adaptation Rule Confidence
from the Rule Generation Process

Vahid Jalali and David Leake

School of Informatics and Computing, Indiana University
Bloomington IN 47408, USA
vjalalib@cs.indiana.edu, leake@cs.indiana.edu

Abstract. Previous case-based reasoning research makes a compelling case for the importance of CBR systems determining the system's confidence in its conclusions, and has developed useful analyses of how characteristics of individual cases and the case base as a whole influence confidence. This paper argues that in systems which perform case adaptation, an important additional indicator for solution confidence is confidence in the adaptations performed. Assessing confidence of adaptation rules may be particularly important when knowledge-light methods are applied to generate adaptations automatically from the case base, giving the opportunity to improve performance by astute rule selection. The paper proposes a new method for calculating rule confidence for automatically-generated adaptation rules for regression tasks, when the rules are generated by the common "difference heuristic" method of comparing pairs of cases in a case base, and a method for confidence-influenced selection of cases to adapt. The method is evaluated in four domains, showing performance gains over baseline methods and case based regression without using confidence knowledge.

1 Introduction

Previous research on CBR confidence, has focused largely on how case and case-base characteristics can be used to estimate confidence (e.g., [1]). An interesting question is how confidence can apply to other CBR knowledge containers to improve confidence estimates for results or even to improve solution quality. For example, for any given level of case confidence, selecting high confidence adaptation rules may improve accuracy. This short paper explores assessing confidence of newly-generated rules, based on the confidence of the data used to generate the rules, and exploiting rule and case confidence information to improve performance.

The paper presents a case study for generating and selecting adaptation rules and selecting cases to adapt for case-based regression tasks, i.e., tasks for which the goal is to generate a numerical value. In the basic form of case-based regression, solutions are generated by k-NN, with values computed by simple averaging approaches. To improve performance, the CBR community has developed a number of knowledge-light methods for generating domain-specific adaptations automatically from the case base. For example Hanney and Keane [2] propose an approach based on applying a difference heuristic to pairs of cases, to generate rules which map similar problem differences to similar solution differences. This paper considers whether it is possible to estimate the

confidence of such rules as they are generated, and how such estimates can affect the performance of case-based regression. It reports on an ablation study which assesses the performance of learning and application with confidence considerations, compared to rule learning and application without, and compared to a baseline of k-NN. It also explores how confidence characteristics of domain cases affect performance of the approach. Experimental results show that using case confidences for selecting base cases and using them in ranking adaptation rules can decrease estimation errors, and that the amount of improvement in each domain varies based on the distribution of the confidence level of the cases.

The paper is organized as follows. Section 2 reviews previous research on solution and adaptation rule confidence estimation. Section 3 introduces our method for assessing the adaptation rule confidence based on the rule generation process. Section 4 shows results of evaluations comparing accuracy of case based regression with and without using confidence information. Section 5 presents conclusions and future work.

2   Previous Research

Solution Confidence Estimation: A number of previous projects have proposed methods for estimating confidence in a CBR system's conclusions by considering characteristics of cases and of the problem space as a whole. For example, Cheetham and Price [3, 4] explore the problem of assigning confidence to solutions in a CBR system by considering similarity scores of the retrieved cases, the deviation of retrieved solutions, etc. [3]. Delany et al. [5] propose estimating classification confidence based on the similarity between the target case and its k nearest neighbors. Reilly et al. [6] propose a feature-based confidence model for assessing confidence of the proposed values for a feature by recommender systems. Mulayim and Arcos [7] propose a method for identifying areas of the problem space for which cases give uncertain solutions, identifying regions of the case base in which those problems are located, to guide maintenance. Hullermeier's [8] Credible Case-based Inference (CCBI), for regression tasks, estimates solutions based on "credible sets" of cases, i.e., sets of high confidence cases. Craw et al. [9] propose using an auxiliary case based reasoning system to predict solution correctness and confidence. Their confidence estimation method works by retrieving a set of adaptation cases with their associated correctness predictions and combining the predictions.

Considering Adaptation Confidence: Distance Weighting: Distance-weighted k-NN can be seen as using a simple proxy for adaptation confidence when solutions are calculated: If confidence in the contributions from different cases depends on their proximity to a query, distance weighted k-NN takes that adaptation confidence into account by weighting nearby cases more heavily.

Determining Confidence using Rule Frequency: Previous research on adaptation rule generation has considered the space of resulting rules, noting that frequency in the pool of generated adaptation rules may give an indication of the reliability of generated rules. Hanney and Keane's [2] seminal work proposes estimating confidence of rules by their frequency. Wilke et al.'s [10] adaptation learning system takes a similar rule-generation approach, estimating rule certainty based on the degree of generalization applied during rule generation.

Provenance-based Confidence Estimation: Leake and Dial [11] propose a provenance-based method for assessing the quality/confidence of adaptation rules by using feedback propagation. Their method assigns blames to applied adaptation rules based on the reported flaws in a solution via feedback. Minor et al. [12] assess the confidence of adaptation results in workflow domains by using introspection for the modified parts of the adapted solutions. For tracking the adaptation process they use provenance information of each used workflow element.

3   Deriving Rule Confidence from the Rule Generation Process

The method introduced in this paper is an extension to the previous work of the authors on an approach called Ensemble of Adaptations for Regression (EAR) [13]. EAR generates adaptation rules by comparing pairs of cases in a local neighborhood around the input problem. Adaptation rules are built by comparing the problem and solution parts of pair of cases and identifying their differences to generate rules that map the observed differences in problems to the observed differences in solutions.

EAR is a lazy approach to adaptation generation. Given an input problem it selects a set of base cases to adapt and also generates a set of adaptation rules as explained above. For adapting the value of each base case it combines the values of the top r rules that most resemble the differences between the base case and the input problem. The final estimation is generated by combining the adapted values of all selected base cases.

EAR selects cases to adapt based on their distance to the input problem and ranks adaptation rules according to the similarity of their problem parts to the corresponding differences between the base cases and the input problem. However, we hypothesize that using confidence knowledge in these two steps can improve the accuracy of the estimations in domains with uncertainty in the values of the cases.

Our approach, which we call confidence-based case-based reasoning (ConfCBR) estimates adaptation rule confidence based on the quality of inputs to the difference heuristic—The confidence of the cases compared to generate the rule. In addition, it uses confidence knowledge in selecting base cases both to adapt cases and for building adaptation rules. Algorithm 1 summarizes the overall process of ConfCBR. In the algorithm, $NeighborhoodSelection\,(Q, n, CB)$ and $RankRules\,(NewRules, C, Q)$ rank base cases and adaptation rules by using (1) and (2) respectively. The subprocesses are described in more detail in the following discussion.

Selecting Cases from which to Generate New Solutions: We hypothesize that the best solutions will be generated by balancing a tradeoff between case similarity and confidence. (This tradeoff is mediated by the quality of case adaptation. If adaptation were always perfect, we would expect the best results always to be obtained by adapting the most confident case.)

Let P be the set of all possible problems, and CB the cases in the case base. Let $distance : P \times P \rightarrow R^+$ measure distance between problem descriptions (for convenience, we will sometimes use the case itself to designate its problem part). Let $conf : CB \rightarrow [0, 1]$ compute case confidence. Then for a case C, ConfCBR calculates the ranking value of that case for base case selection by:

Algorithm 1 ConfCBR's method of estimating target values

Input:
Q: input query
n: number of base cases to be used
CB: case base
Output: Estimated solution value for Q

    $CasesToAdapt \leftarrow$ NeighborhoodSelection(Q,n,CB)
    $NewRules \leftarrow$ RuleGenerationStrategy(Q, CB)
    for C in $CasesToAdapt$ do
        $RankedRules \leftarrow RankRules\,(NewRules, C, Q)$
        $ValueEstimate_C \leftarrow CombineAdaptations\,(RankedRules, C)$
    end for
    return CombineVals($\cup_{C \in CasesToAdapt} ValEstimate_C$)

$$rank\,(C, Q) \equiv \frac{conf\,(C)^{\alpha}}{distance\,(C, Q)} \tag{1}$$

where α is a positive real number whose values tune the base case rankings for different domains. If α is equal to zero, cases will be ranked merely based on their distance to the input query. Increasing the value of α has the effect of assigning higher rankings to more confident cases, with large values of α asymptotically approaching assigning equal ranking values (i.e. zero) to cases that are not 100% confident.

Adaptation Rule Ranking: To rank candidate adaptation rules generated from the case base by the difference heuristic, ConfCBR computes a ranking score based on two factors: (1) confidence of the cases from which the rules were generated, and (2) how close the cases from which the rule was generated are to the case to adapt. More specifically, let $R_{i,j}$ be the adaptation rule built from C_i and C_j, and $\Delta\,(C, Q)$ be the difference vector of the features of the query Q. The ranking value of $R_{i,j}$ for adapting the solution of C is:

$$rank\,(R_{i,j}, C, Q) \equiv \frac{(conf\,(C_i) \times conf\,(C_j))^{\beta}}{distance\,(Q_{i,j}, \Delta\,(C, Q))} \tag{2}$$

where β is a positive real number whose values tune the ranking of adaptation rules in different domains.

4   Evaluation

Our experiments address the following questions about ConfCBR:

1. How does ConfCBR's accuracy compare to its accuracy in the ablated conditions (1) rule confidence considered, case confidence ignored, (2) case confidence considered, rule confidence ignored, (3) both confidence factors ignored.

2. How does ConfCBR's accuracy (using confidence, and without) compare to the baseline of distance-weighted k-NN?
3. How do varying case confidence distributions in the case base affect ConfCBR's accuracy?

4.1   Experimental Design

Experiments applied ConfCBR for case-based regression in four sample domains from the UCI repository [14]: MPG, Auto, Hardware and Housing. A data cleaning process removed cases with unknown values and discarded the symbolic features. For each feature, values were standardized by subtracting that feature value's mean from each individual feature value and dividing the result by the standard deviation of that feature. Leave-one-out testing is used for all domains and estimation errors are calculated in terms of Root Mean Squared Error (RMSE). Hill climbing was used to determine the values of α and β for each domain.

Adaptation rules are generated by comparing the top 5% nearest neighbors of the input query, using Euclidean distance in (1) and (2). The top 5% cases are ranked and selected by using (1). These cases are used both as base cases to adapt and the source cases for generating adaptation rules. The case-based regression system then ranks adaptation rules by using (2) and applies a set of rules for adapting the solution of each base case. The number of adaptations applied per base case is also determined by a hill climbing process in all domains. The final estimates are generated by combining adapted solutions of the selected base cases.

The goal of our current study is not to generate case confidence values, but rather, to assess how confidence information can be exploited, once it has been generated. To evaluate our approach under controlled conditions for which the quality of confidence estimations is known, we generated test data whose correctness was characterized by varying known confidence values, as follows. First, we randomly assigned confidence levels to the cases by a Gaussian distribution, with 0.8 and 0.2 used as the mean and standard deviation of the confidence level distributions in all domains except explicitly stated otherwise. The stored values of the cases were then adjusted randomly, according to the assigned confidence values. For example, if 0.9 is assigned to a case as its confidence value, its stored value is increased or decreased by 10% of its value. The original value of the case is used as the "correct" value for assessing performance.

4.2   Performance Comparison

The Effect of Using Confidence Knowledge   Fig. 1 depicts the RMSE for CBR regression (without using confidence knowledge), using confidence knowledge in selecting the base cases only (ConfCBRC), using confidence knowledge only for ranking adaptations only (ConfCBRR) and using confidence knowledge for both selecting base cases and ranking adaptation rules (ConfCBR), for each test domain. As expected, in all domains CBR without confidence considerations shows the worst performance. In three of the four test domains (all except MPG) using confidence knowledge in ranking adaptation rules is more successful in decreasing estimation error compared to using

confidence knowledge for selecting the base cases to adapt. In almost all cases, using case confidence knowledge both for selecting base cases and ranking adaptation rules (ConfCBR) provides the most accurate results. Exceptions occur only for one configuration of the Auto domain (when solutions are generated from 5 base cases) and when 5 or more base cases were used in the Hardware domain. In both those cases, confidence-based rule ranking only outperformed the combination, but the combination outperformed rule ranking only for smaller numbers of base cases.
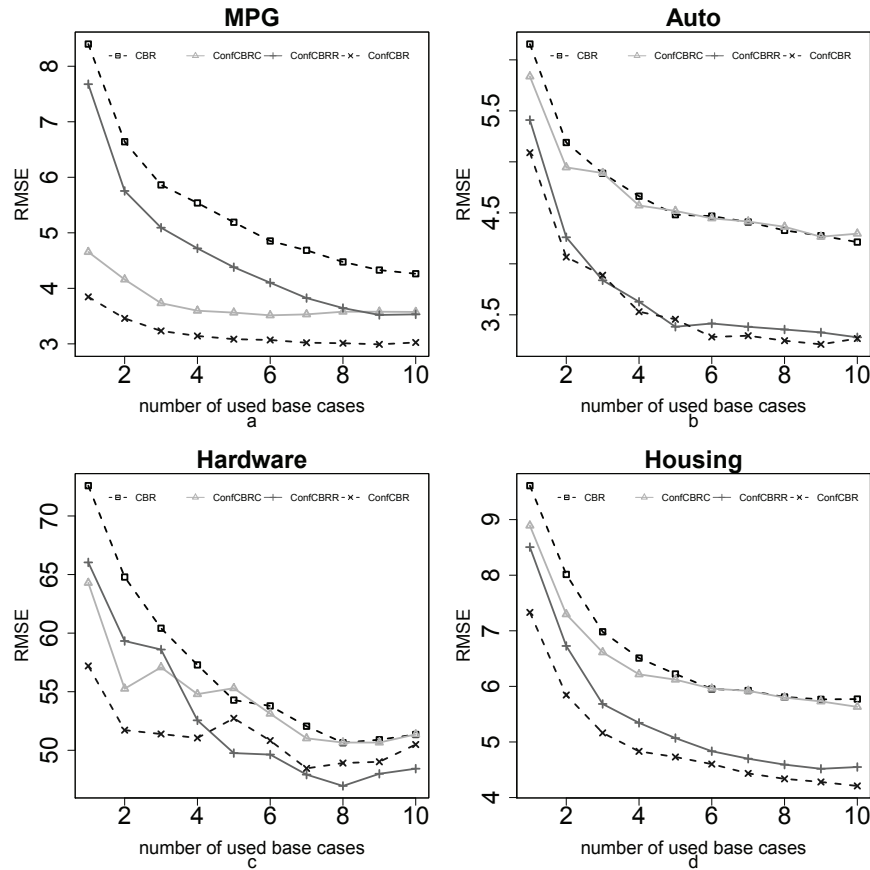

Fig. 1. RMSE comparison for no use of confidence (CBR) only using confidence for ranking base cases (ConfCBRC), only using confidence for ranking adaptation rules (ConfCBRR) and using confidence for both (ConfCBR)

ConfCBR vs. k-NN To compare the accuracy of ConfCBR with a baseline, we conducted experiments in the test domains using conventional distance weighted k-NN,

and distance-weighted k-NN enhanced with case confidence knowledge (ConfkNN). The confidence knowledge in ConfkNN is used for selecting the cases from which the solution will be generated by using (1).

Fig. 2 shows the RMSE of k-NN, Confidence based k-NN (ConfkNN), CBR and ConfCBR in the test domains. For three out of four domains (all except Hardware), the worst performance belongs to the basic CBR approach, which reflects its inability to adjust to varying confidence levels (either of base cases or the cases from which adaptations are built). The largest performance gap between k-NN and ConfCBR is observed for the Hardware domain (ConfCBR performs 33% better than ConfkNN) while this gap is minimized in the MPG domain (ConfCBR only performs 3% better than ConfkNN). In all domains ConfCBR performs better than the baseline methods, showing that CBR enhanced with confidence knowledge is able to generate more accurate estimations compared to the other tested alternative methods.
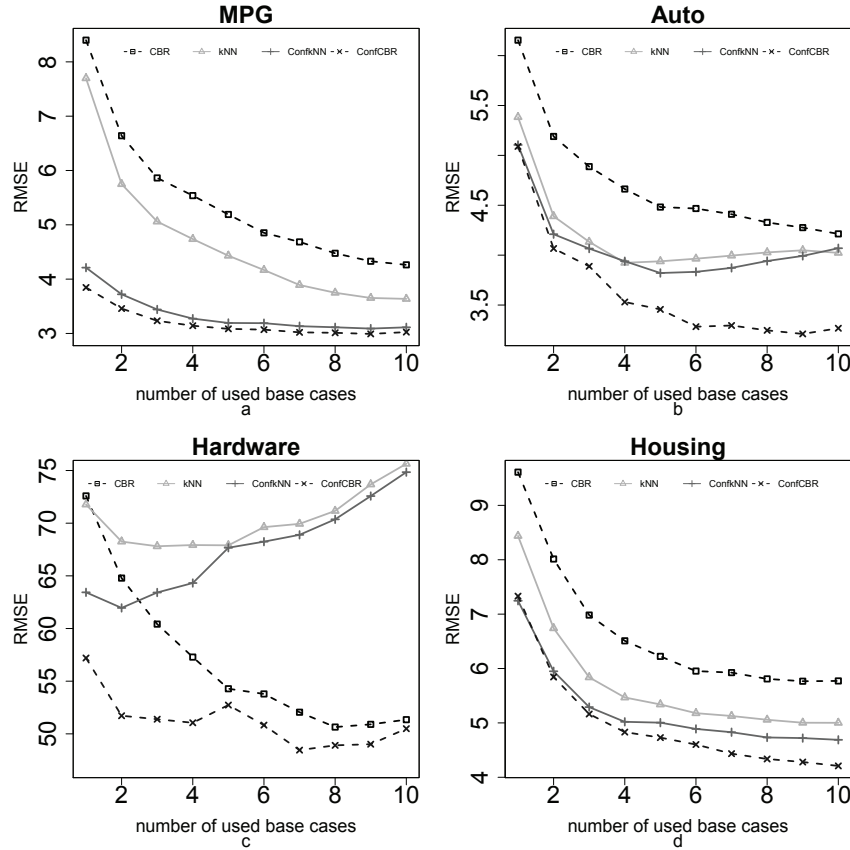

Fig. 2. RMSE of CBR, ConfCBR, k-NN, and k-NN enhanced with confidence knowledge (ConfkNN)

The Effect of Case Confidence Level Distribution on ConfCBR To assess the effect of case confidence level distribution on ConfCBR we conducted experiments in the Housing domain with different case confidence level distributions. Fig. 3 shows the RMSE of CBR and ConfCBR in the housing domain for four different confidence distributions. As a reference, Part a of Fig. 3 repeats the results of Part a of Fig. 1. However, parts b, c and d show results for three new distributions. Based on comparison of parts a, b and c of Fig. 3 we hypothesize that the difference between the relative performance of CBR and ConfCBR depends more on the standard deviation of the case confidence levels than on their mean value. Part d suggests that for relatively small standard deviations, performance of CBR and ConfCBR is almost identical.
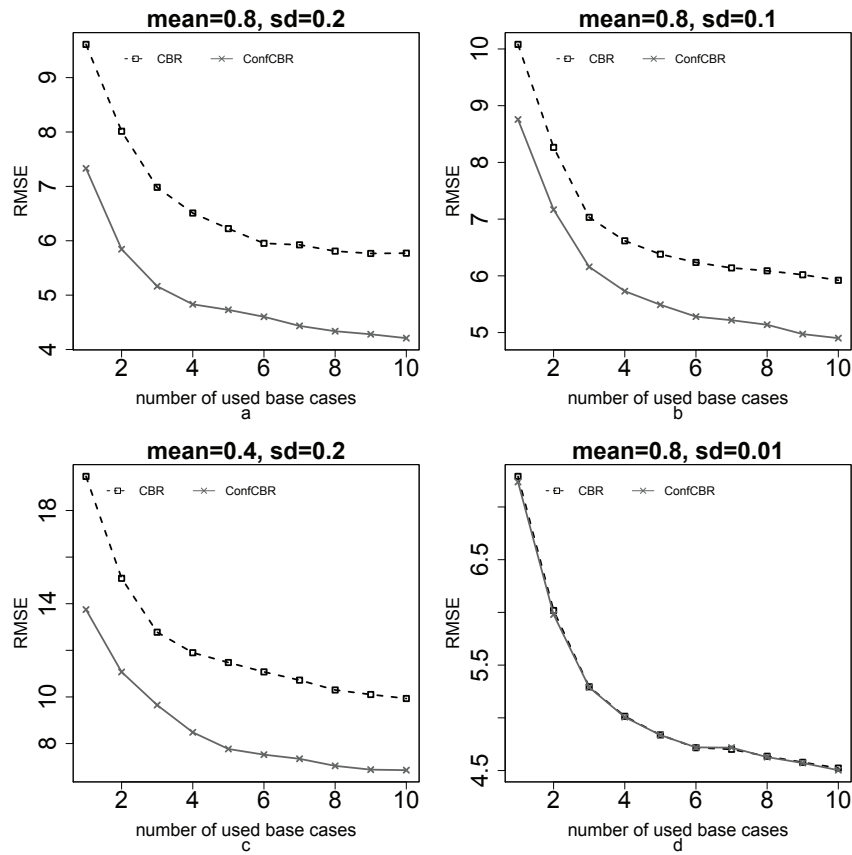

Fig. 3. RMSE of CBR without using confidence knowledge (CBR) and CBR using case confidence knowledge both for ranking base cases and adaptation rules (ConfCBR) in the Housing domain for four sample normal distributions of case confidence levels.

5   Conclusion and Future Work

This paper has explored how considerations of case confidence can help to assess the confidence of adaptation rules generated by the difference heuristic, and how adaptation rule and case confidence can be brought to bear not only to assess confidence of solutions, but to generate better solutions. It has introduced a new method, confidence-based case based regression (ConfCBR), which uses confidence knowledge both for selecting base cases and ranking adaptations. Experimental results showed that ConfCBR outperforms a corresponding case-based approach to regression without confidence knowledge and k-NN baseline methods in four sample domains, often by substantial margins. Results also showed that the benefit depends significantly on the distribution of case confidence levels in a case base.

Our current investigation was based on sample data for which confidence levels were artificially created. We are developing methods for estimating the case confidences of new case bases by statistical methods such as outlier detection, and intend to examine ConfCBR performance when case confidence levels are estimated automatically. Another future direction is studying other methods for estimating rule confidence.

References

1. Cheetham, W., Price, J.: Measures of solution accuracy in case-based reasoning systems. In: Proceedings of the Seventh European Conference On Case-Based Reasoning. (2004) 106–118
2. Hanney, K., Keane, M.: The adaptation knowledge bottleneck: How to ease it by learning from cases. In: Proceedings of the Second International Conference on Case-Based Reasoning, (1997)
3. Cheetham, W.: Case-based reasoning with confidence. In: Fifth Euorpean Workshop on Case-Based Reasoning. (2000) 15–25
4. Cheetham, W., Price, J.: Measures of solution accuracy in case-based reasoning systems. eccbr 2004. In Funk, P., González, P., eds.: ECCBR-2004: Advances in Case-Based Reasoning, (2004) 106–118
5. Delany, S.J., Cunningham, P., Doyle, D., Zamolotskikh, A.: Generating estimates of classification confidence for a case-based spam filter. In: Proceedings of the 6th international conference on Case-Based Reasoning Research and Development. ICCBR'05, (2005) 177–190
6. Reilly, J., Smyth, B., McGinty, L., McCarthy, K.: Critiquing with confidence. In: ICCBR. (2005) 436–450
7. Mülâyim, O., Arcos, J.L.: Understanding dubious future problems. In: Proceedings of the 9th European conference on Advances in Case-Based Reasoning. ECCBR '08, (2008) 385–399
8. Hullermeier, E.: Credible case-based inference using similarity profiles. IEEE Trans. on Knowl. and Data Eng. **19**(6) (June 2007) 847–858
9. Craw, S., Jarmulak, J., Rowe, R.: Learning and applying case-based adaptation knowledge. In Aha, D., Watson, I., eds.: Proceedings of the Fourth International Conference on Case-Based Reasoning, (2001) 131–145
10. Wilke, W., Vollrath, I., Althoff, K.D., Bergmann, R.: A framework for learning adaptation knowledge based on knowledge light approaches. In: Proceedings of the Fifth German Workshop on Case-Based Reasoning. (1997) 235–242

11. Leake, D., Dial, S.: Using case provenance to propagate feedback to cases and adaptations. In: Proceedings of the Ninth European Conference on Case-Based Reasoning, Springer (2008) 255–268
12. Minor, M., Islam, S., Schumacher, P.: Confidence in workflow adaptation. In: ICCBR. (2012) 255–268
13. Jalali, V., Leake, D.: Extending case adaptation with automatically-generated ensembles of adaptation rules. In: ICCBR. (2013)
14. Frank, A., Asuncion, A.: UCI machine learning repository (2010) http://archive.ics.uci.edu/ml.