Experiments in Curvature-Based Segmentation of Range Data

Emanuele Trucco and Robert B. Fisher
Department of Artificial Intelligence,
University of Edinburgh, Edinburgh, Scotland.

Abstract

This paper focuses on the experimental evaluation of a range image segmentation system which partitions range data into homogeneous surface patches using estimates of the sign of the mean and Gaussian curvatures. We report the results of an extensive testing programme aimed at investigating the behavior of important experimental parameters such as the probability of correct classification and the accuracy of curvature estimates, measured over variations of significant segmentation variables. Evaluation methods in computer vision are often unstructured and subjective: this paper contributes a useful example of extensive experimental assessment of surface-based range segmentation.

1 Introduction

"In all the attempts made to date to show that $2 + 2 = 4$, no one has ever considered the windspeed."
R. Queneau, *Quelques remarques sommaires relatives aux propriétés aérodynamíques de l'addition.*

In this paper we present an example of experimental performance assessment for a range surface segmentation technique and its implementation. The results of the performance tests provide very useful experimental information for choosing consistent system parameters and thresholds as well as for deriving heuristics. The paper suggests also a testing scheme implicitly, by identifying relevant parameters and investigating their behavior experimentally. Moreover, quantitative assessment allows comparison with implementations addressing the same problem.

We give here only a very succinct account of our segmentation technique for reasons of space; the reader is referred to [11] and [10]. Our approach adopts Gaussian smoothing, modeled within the *diffusion* paradigm. The longer the diffusion interval, the larger the equivalent Gaussian kernel. Depth and orientation discontinuities maps are precomputed and used to restrict the diffusion process to non-discontinuity points, thus avoiding the creation of spurious curved regions around discontinuity contours. In order to cope with the distortion introduced near surface boundaries, we enforce a shape-preserving adaptive boundary condition [10]. Then the program computes a *qualitative description* of the range surfaces using estimates of the sign of the mean and Gaussian curvatures (H and K respectively). Morphology is used to improve the quality of the raw HK maps. Planar points are detected by thresholding the H,K maps. The zero threshold on Gaussian curvature, K_0, is computed *consistently* (in the sense that small perturbations in the principal curvatures will still lead to correct estimates of H and K signs) from that on mean curvature, H_0, thus eliminating one user-defined parameter. The final surface description consists of a number of *patches*, formed by surface points of the same qualitative shape ($K = H = 0$ planar, $K < 0$ hyperbolic, $K > 0$ elliptic, $K = 0 and H \neq 0$ cylindrical; $sgn(H)$ gives convexity for the last two cases). The surface description is qualitative in the sense that it uses only the shape class to identify and characterize the patches.

An example of segmentation of a real image, the well-known Renault part (size about 190x100x80mm^3) is shown in Figure 1. This is a complex sculptured object with both developable and non-developable surfaces. The H threshold was experimentally set at $H_0 = 0.024$. Consistent thresholding yields $K_0 = 0.00422$. The image was smoothed between $t_0 = 0$ and $t_{max} = 9$ (diffusion time interval), equivalent to $\sigma_{max} = 3$ of the equivalent Gaussian kernel. The result of the segmentation is shown in Figure 2. Most of the significant patches appear in the final segmentation, providing a significant description for model-based tasks. Our segmentation compares well with most results reported in the literature for the same object. For instance, Besl and Jain's segmentation[1], obtained in comparable experimental conditions (128x128 pixel, 8-bit depth, very similar pose), looks much more fragmented than our result. Results of techniques for planar patch segmentation (e.g. [5]), are somewhat more difficult to compare with ours, owing to the different nature of the description computed.

We have identified several parameters that we consider important for assessing the performance of a HK segmentation system. These are the probability of correct classification, the accuracy of mean curvature estimates, the distortion on planar patches, the amount of shape distortion near patch boundaries and the sensitivity of curvature measurements to size (small patches). All these were measured in a large number of tests run with synthetic patches, with quantized and double precision data, against systematic variations of many experimental variables, including shape class, amount of smoothing, patch orientation, patch size, and zero curvature thresholds. Extrema values for the ranges of the independent variables were determined in advance by trial and error. The quantization step of the independent variables was kept as low as possible in order to guarantee significant resolution. The program's performance was also analyzed with a number of synthetic and real images. We used synthetic planar, cylindrical, spherical and saddle patches of varying shape and size. All images were 128x128 pixel, with quantized or double precision values. Quantized data were obtained by rounding off depth values to the nearest integer, thus forcing a quantization of step $\Delta h = 1$ over a range of 256 possible values. Real range data were acquired using the laser striper equipment developed in our laboratory (accuracy of about one part in 10,000).

Among the many works appeared recently on range image segmentation (e.g. [1], [8], [7], [9], [4]), Flynn and Jain report a most interesting empirical analysis of the accuracy of five curvature estimation techniques [6]. They used synthetic patches of various shape classes (but not hyperbolic). Their tests were run over narrower and more coarsely sampled parameter ranges than ours, for instance using radii of 1, 2, 5 and 10 units only. No results on classification and sensitivity were reported. All conclusions, wherever applicable, agree with ours. The overall common message is that curvature-based surface descriptions are interesting for range vision, but qualitative methods seem more reliable than quantitative ones. Yang [12] adopted structured lighting (grids) to measure the principal curvatures at stripe junctions with several synthetic shapes and real images. In all cases he used large patches, reporting error variances between 6% and 23%. He noticed some of the problems we have investigated (e.g. the need for large patches to provide reliable curvature estimates) but did not perform any specific analysis. Several surface-based methods have been directed at the detection of planar surfaces. For instance, Bahnu [2] segmented multiple range images of the Renault part, acquired from different viewpoints, to create a 3-D surface model. The surface was then divided into planar patches. Surface patch detection and classification algorithms with planar and curved patches have been reported by various authors. Cai [3] analyzed the merits and limits of HK segmentation in scale-space, identifying stable features that can be reliably traced (for instance C^2 discontinuities), and introduced the consistent zero thresholding technique that we have used in our system.

2 Testing probability of correct classification

Classification tests were run by measuring the percentage of correctly classified points in synthetic patches against variations of the patch geometry (shape class and curvature), the number of smoothing cycles and the zero threshold H_0 (the only user-defined zero threshold for curvature, K_0 being obtained from H_0 by consistent thresholding). An optimal H_0 range was calculated for each patch by estimating experimentally the values generating 100% and 0% percentages (or sufficiently high and low percentages if the 0%-100% H_0 range was too large). This range was then divided into 20 intervals to give the H_0 step for the given patch. For the nc range, we found experimentally that $nc = 1, 2, 3$ provided enough smoothing without introducing unacceptable distortions. Similar values, expressed as standard deviation of Gaussian kernels, were used in Flynn and Jain's experiments [6].

Figure 3 shows the graphs obtained for cylinders, with ideal (double precision) data and noisy (quantized) data. The results reveal a high sensitiveness: classification varies from 50% to 90% rather quickly. Sensitiveness decreases for $nc = 3$. Moreover, reliable H_0 values correspond to overestimated radii, about three times the true values. This is due to the averaging effect of Gaussian smoothing. For radii up to 25 pixels $H_0 = 33 \cdot 10^{-3}$ is a good lower bound; this works admissibly for planes of up to 80° normal slope. Moreover, examples of useful heuristics suggested by the results are: use $H_0 = \frac{1}{2} \max H$ of object expected; do not use $H_0 \leq 0.015$ for $nc = 3$, $H_0 \leq 0.027$ for $nc = 2$. For spheres and cylinders the values of the radius equivalent to the H_0 value yielding the given percentage of correct classification was plotted against the patch radius. The corresponding H_0 was obtained as $H_{cyl} = \frac{1}{2R_{cyl}}$ and $H_{sph} = \frac{1}{R_{sph}}$. For saddles and planes, H_0 was plotted against the geometric parameter of the patch (e.g. the elevation of the normal from $z = 0$ for planes, the eccentricity for saddles).

The patch classification graphs suggest also that the HK thresholding mechanism is inadequate for fine distinction between planes and low-curvature curved surfaces. Figure 3 indicates that worst-case misclassifications might occur with cylinders of radius about 18 pixels or spheres of radius about 35 (corresponding to the worst distortion in Figure 3). However, for a large interval of plane slopes, approximate interference values are $R = 70$ for cylinders and $R = 140$ for spheres. The more accurate the classification (i.e. the more points correctly classified), the higher the equivalent curvature radius, because less and less points are classified as planar regions on curved patches as the threshold H_0 decreases. Notice that H_0 values necessary to achieve good classification percentages are always smaller than the real H for the patch, an expected effect of Gaussian smoothing. The curves obtained from quantized patches are more irregular, because the quantization introduces shape discontinuities on smooth surfaces, which results in H estimates significantly different from the ideal values. In general, the better the classification percentage the more irregular the graph, because points get included whose H values are farther and farther away from the ideal values.

3 Testing accuracy of mean curvature estimation

The accuracy of H estimates is severely limited by the combined effects of *quantization* and *smoothing*. The quantization noise can be attenuated by Gaussian smoothing at the cost of altering the shape of the surface. Numerical errors introduced by the computation play a much less important role if compared with the previous two factors and are not investigated in the following. We discuss briefly two sets of experiments aimed at assessing the quality of numerical curvature estimates, using synthetic and real images of planes and cylinders.

We ran tests to measure H on cylindrical and spherical surfaces of different sizes with increasing amounts of smoothing ($nc = 1, 2, 3$). For reasons of space we show only one graph synthesizing the results in Figure 4, which plots the mean and standard deviation of H for cylindrical patches of increasing radius, smoothed with two diffusion intervals ($nc = 2, 3$). The estimates get worse as the diffusion interval increases, as expected. For a sphere of radius $r = 19$, for instance, the deviation from the expected value is 0.0026, 0.0037 and 0.0054 for 1, 2 and 3 cycles of smoothing respectively; the correspondent percentage error in the estimated radius is 5%, 7% and 11% respectively. The estimates become more and more unreliable as the radii decrease. The results suggest that, for radii less than 10 pixels, the estimates degrade rapidly as the number of smoothing cycles increase; again in the sphere case, the error in the estimated radius for $r = 10$ is already 31% for one smoothing cycle, and becomes 66% and 167% for two and three smoothing cycles. The reason is, as the size of the smoothing Gaussian approaches that of the surface to be smoothed, the curvature of the latter is more and more distorted. The practical importance of empirical distortion estimates is that they could be used for correcting the estimates of H (and similarly for the principal curvatures) with real surface data.

4 Testing distortion

Plane distortion. In order to establish the exact limits of curvature thresholding in our implementation, we measured the mean curvature on variously oriented planar patches against varying slopes and smoothing intervals. The results indicate that quantized values always lead to worse distortions. With $nc = 1$ the quantization noise is still rather strong; increasing the value of nc attenuates the noise better but makes the distortion introduced by smoothing worse. For $nc = 2, 3$ the latter effect is not yet too evident; however, the worst standard deviation measured on the planes was about 0.05 ($nc = 1$, normal elevation 80°, plane slope 20°), corresponding to a spherical curvature radius of 20 pixel. The corresponding mean curvature was about $65\text{x}10^{-5}$, corresponding to a curvature radius of $153\text{x}10^{3}$ pixel: in practice a plane. In general, the distortion in the mean values are negligible, with very large equivalent spherical radii; but the standard deviations are three orders of magnitude larger, indicating the presence of a serious number of outliers.

Distortion at boundaries. The boundary condition method we adopted to contrast shape distortion at boundaries is called *adaptive leakage* [10]. Instead of "sealing" the surface to be smoothed, or imposing a fixed interaction between surface and background during diffusion smoothing (*fixed leakage* [3]), our method adapts to the local conditions to minimise smoothing-induced distortions. We compared the performance of fixed and adaptive leakage by smoothing a number of quantized planar patches at different orientations and observing the distortion arising. Distortions concentrated along patch boundaries, generating long, slender stripes of misclassified pixels. We recorded the percentages of misclassified pixels for each orientation of the planar patch and averaged over the various orientations. The results indicate that adaptive leakage reduces the percentages by at least one order of magnitude [11], keeping the noisy patches created by boundary distortion effects very small. Such patches can be eliminated by the subsequent morphological enhancement.

5 Patch shrinking

Shrinking is introduced by different causes, which include the fact that pixels belonging to region boundaries are not included in the regions (because of the ambiguity regarding which surface they belong to), discontinuity detection (contours will be shrunk by an amount which depends on the local surface slope) and morphological HK image enhancement. The worst case shrinkage in pixels introduced after discontinuity detection and smoothing is $(1 + D)_{disc} + 1_{bound}$, where 1-pixel shrinkages are contributed by discontinuity marking and boundary treatment, and D is the additional shrinking caused by discontinuity detection at sloping surface borders. It is difficult to estimate the size of the smallest patch detectable. This varies with the type of patch, its orientation, the number of smoothing cycles, and the effect of morphological transformations on the HK sign images. An obvious limit is imposed by the shrinking discussed above. Using the results of our tests we arrived at a size threshold of about 14 pixels diameter for *reliable* patch classification. Below that value spherical and cylindrical patches are significantly distorted by amounts of smoothing adopted normally. Hyperbolic patches are more stable, the only trouble being possibly that low-curvature patches might still be misclassified as planes.

6 Conclusions

Perhaps the most important practical lesson learnt from our performance testing programme is that curvature estimates are *very* sensitive to quantization noise. Smoothing is required to get stable curvature estimates, but at the price of introducing a general reduction in the estimated curvatures. As a consequence, H_0 values ensuring good classification percentages are always smaller than the expected H values. Gaussian smoothing and quantization noise both contribute to distorting small features. Our system achieved really reliable sign classification with patches larger than about 13 pixel side (referring to 128x128 images smoothed with $nc \leq 3$). The accuracy of H magnitude estimates is also limited. Even with large patches, accuracies better than about 10% are difficult to achieve. In the worst cases, estimates of $|H|$ can be off by up to 200% in our tests. Considering our results together with those obtained by Flynn and Jain [6], it is apparent that the sign of the mean and Gaussian curvatures can be computed more reliably than curvature magnitude. Quantization and smoothing make it also difficult to distinguish accurately between planar patches and curved surfaces with low curvatures. Interference radii in our experiments are about $R = 70$ for cylinders and $R = 140$ for spheres. Whether this is satisfactory or not depends ultimately on the particular application.

A number of useful indications emerged from the tests, and we mention some here. Given the fragility of the HK technique as plane detector, the task of segmenting a scene composed of both planar and curved surfaces could be best solved by running a robust, specialized plane fitting techniques first, in order to single out the planar patches (as tried by Yang [12]); the remaining curved surfaces could be then segmented with an appropriately low H_0. Several heuristics can be inferred from graphs of the kind presented in this paper. For example, with smooth patches, a reasonable value for H_0 is $H_0 = \frac{1}{2}H_0^{max}$, where H_0^{max} is the maximum expected mean curvature on the patch considered. Another example is that the use of H_0 less than 0.015 with $nc = 2$ and less than 0.027 for $nc = 3$ can lead to unreliable results. Tabular corrections based on H measurements might be possible to contrast empirically the distortion introduced by smoothing, at least in controlled conditions.

The need for experimental verification of algorithms and well-defined testing criteria has been pointed out repeatedly in computer vision. In the framework of a discipline getting richer and richer in tools but remaining significantly poor in evaluation methods, we feel that this paper gives several valuable contributions. First, the tests reported suggest implicitly guidelines for extensive experimental assessment of surface-based range segmentation. Second, the results provide practical quantitative measurements and indications about a popular curvature-based segmentation technique, useful for example as an empirical basis for investigating possible corrections for curvature distortion caused by smoothing. Finally, the results can be used to select consistent thresholds for parameters typically adopted in range segmentation.

References

[1] P. Besl and R. Jain: "Segmentation through Variable-Order Surface Fitting", *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-10, pp. 167 – 192, 1988.

[2] B. Bhanu, "Representation and Shape Matching of 3-D Objects", *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, pp. 340 – 350, 1984.

[3] L. D. Cai, "A 'Small Leakage' Model for Diffusion Smoothing of Range Data", *Proc. 11th IJCAI*, Detroit, pp.1585–1590, 1989.

[4] T.-J. Fan, G. Medioni and R. Nevatia, "Segmented Descriptions of 3-D Objects", *IEEE Trans. Robotics Automat.*, vol. RA-3, 1987.

[5] O.D. Faugeras, M. Hebert and E. Pauchon, "Segmentation of Range Data into Planar and Quadratic Surfaces", *Proc. IEEE CVPR Int. Conf. Comp. Vision Pattern Rec.*, pp. 8 – 13, 1983.

[6] P.J. Flynn and A.K. Jain, "On Reliable Curvature Estimation", *Proc. IEEE CVPR Int. Conf. Comp. Vision Pattern Rec.*, pp. 110 – 116, 1989.

[7] R. Hoffman and A. Jain, "Segmentation and Classification of Range Images", *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, no. 5, pp. 608 – 620, 1987.

[8] D.J. Ittner and A.K. Jain, "3-D Surface Discrimination from Local Curvature Measures", *Proc. IEEE CVPR Int. Conf. Comp. Vision Pattern Rec.*, pp. 119 – 123, 1985.

[9] J. Ponce and M. Brady, "Towards a Surface Primal Sketch", *Proc. IEEE Int. Conf. on Robotics and Automation*, St.Louis, MO, March 1985.

[10] E. Trucco, "On Shape-Preserving Boundary Conditions for Diffusion Smoothing", *Proc. IEEE Int. Conf. Robotics and Automation*, Nice, pp. 1690–1694,1992.

[11] E. Trucco and R. B. Fisher, "Computing Surface-Based Representations from Range Images", *Proc. IEEE Int. Symp. on Intelligent Control*, Glasgow, pp. 275 – 280, 1992.

[12] H.S. Yang, "Range Image Segmentation and Classification Via Split-and-Merge on Surface Curvatures", *Proc. 4th Int. Conf. on Pattern Recogn.*,British Pattern Rec.Soc., Cambridge, pp. 58–67, 1988.

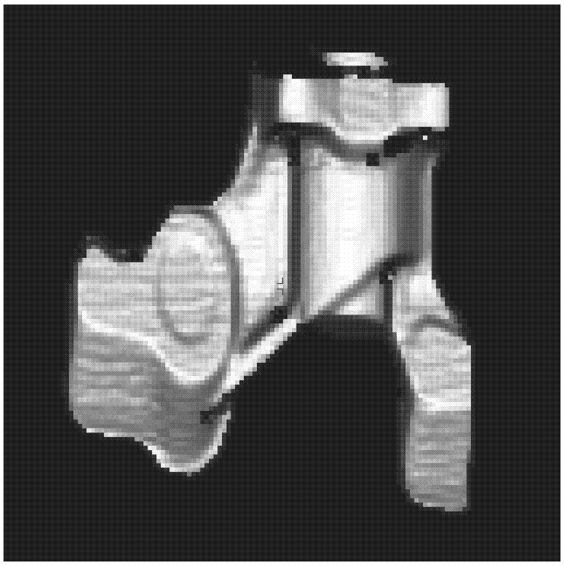

n=2

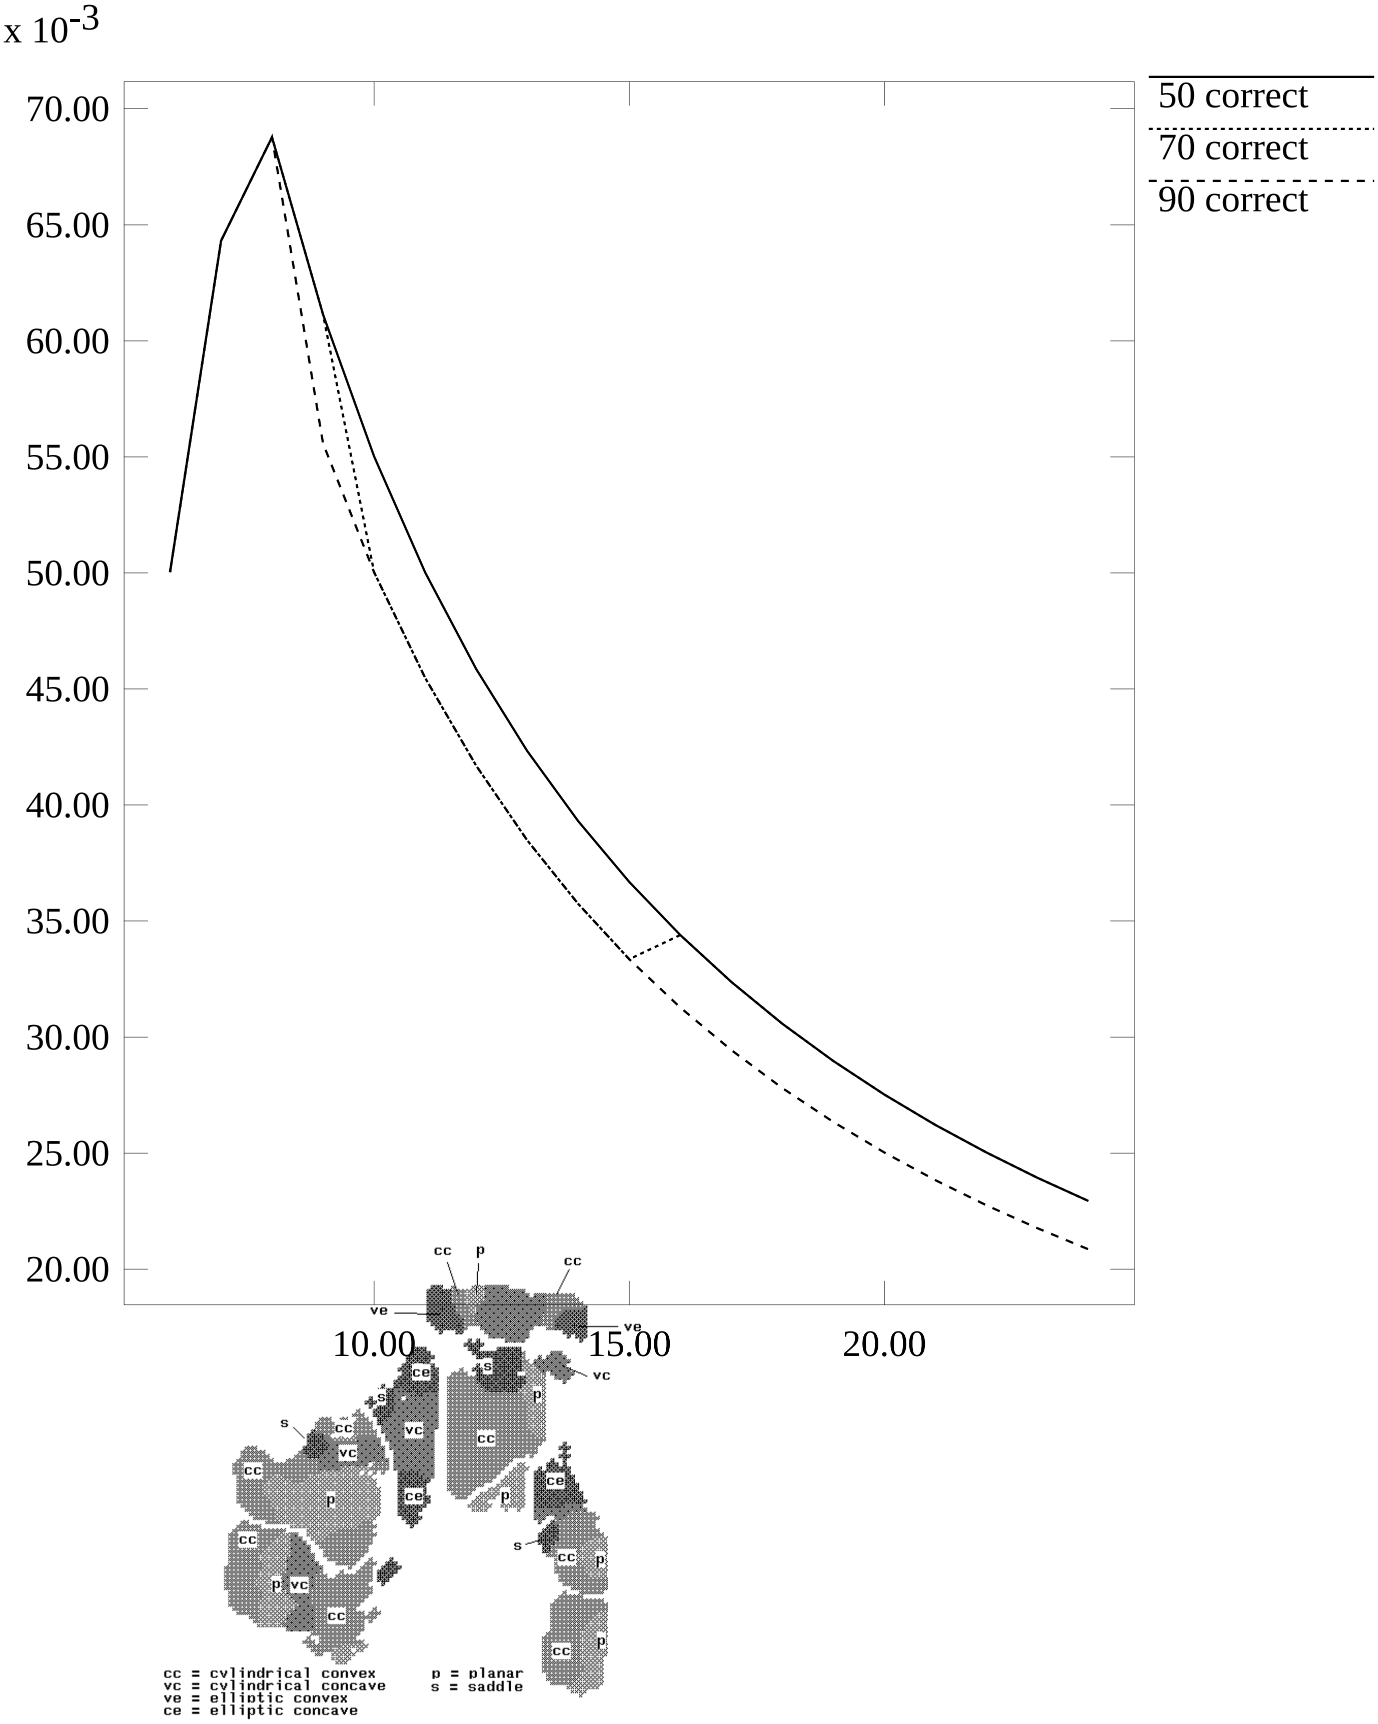
x 10^-3
70.00
65.00
60.00
55.00
50.00
45.00
40.00
35.00
30.00
25.00
20.00
10.00
15.00
20.00
50 correct
70 correct
90 correct
cc
p
ve
s
vc
ce
cc = cylindrical convex
vc = cylindrical concave
ve = elliptic convex
ce = elliptic concave
p = planar
s = saddle

n=2

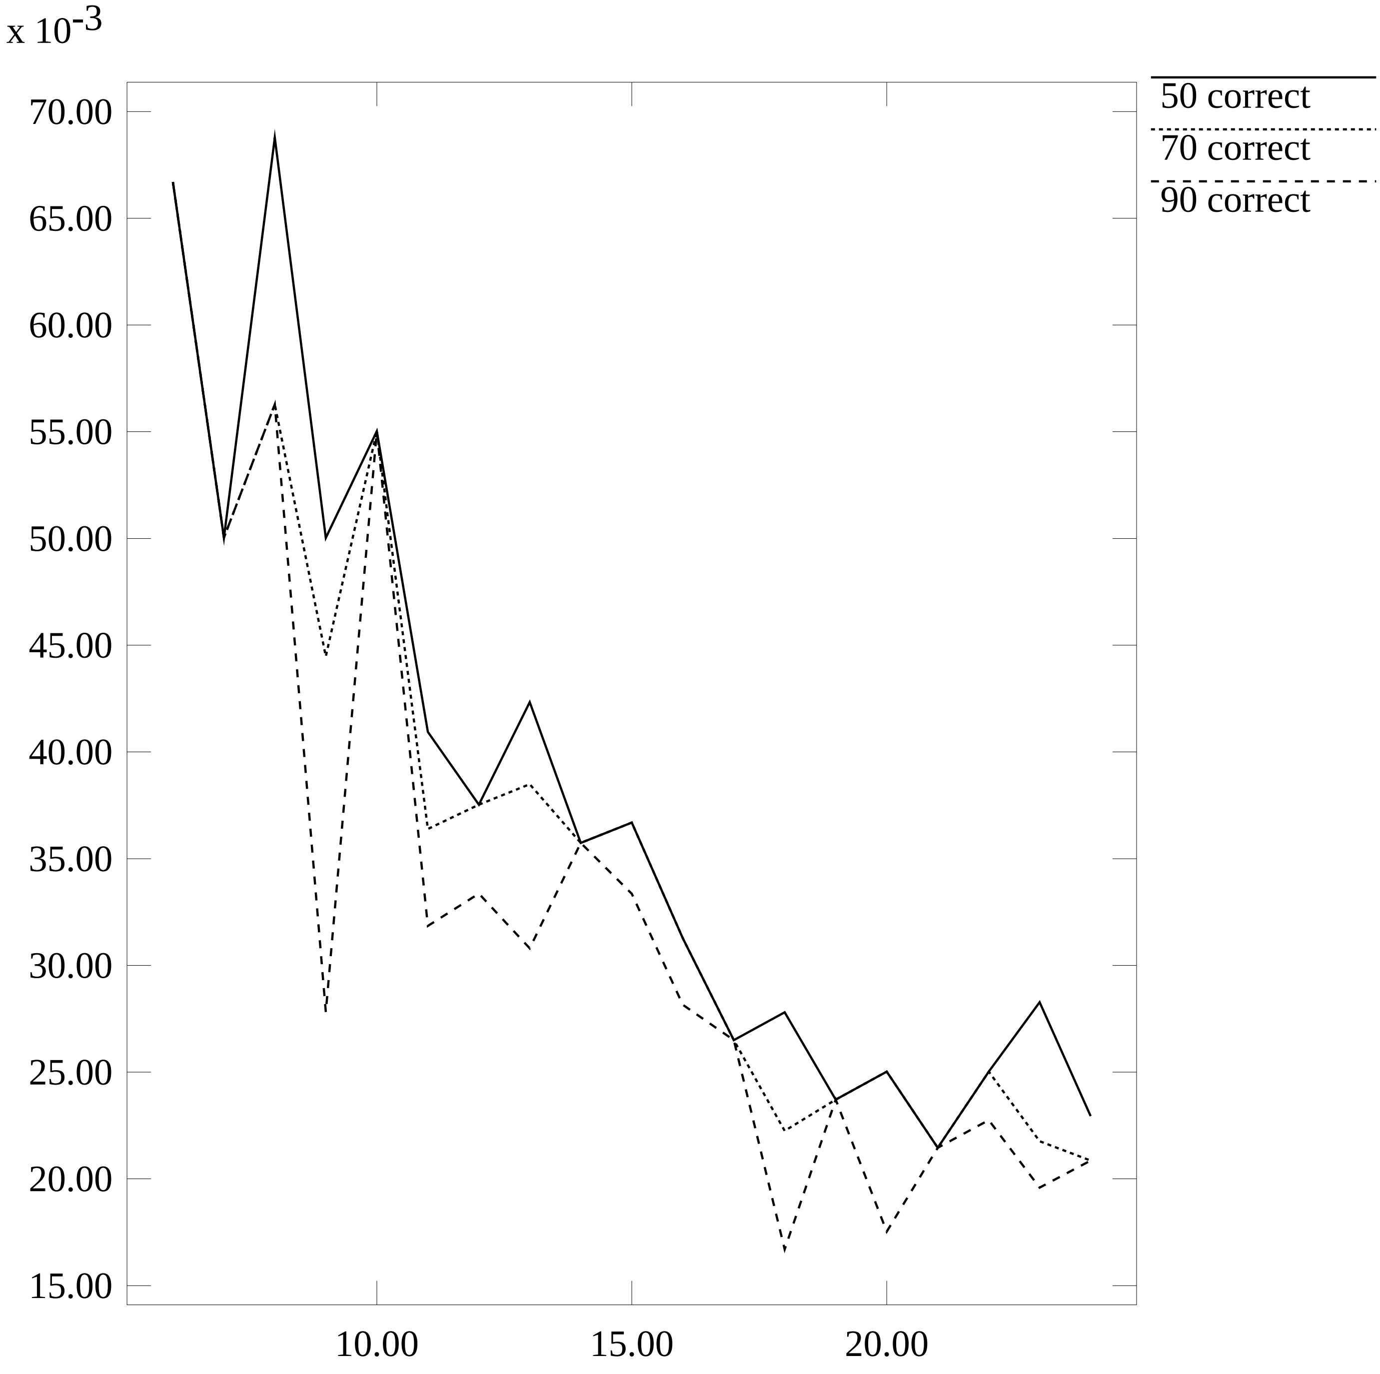
x 10-3
70.00
65.00
60.00
55.00
50.00
45.00
40.00
35.00
30.00
25.00
20.00
15.00
10.00
15.00
20.00
50 correct
70 correct
90 correct

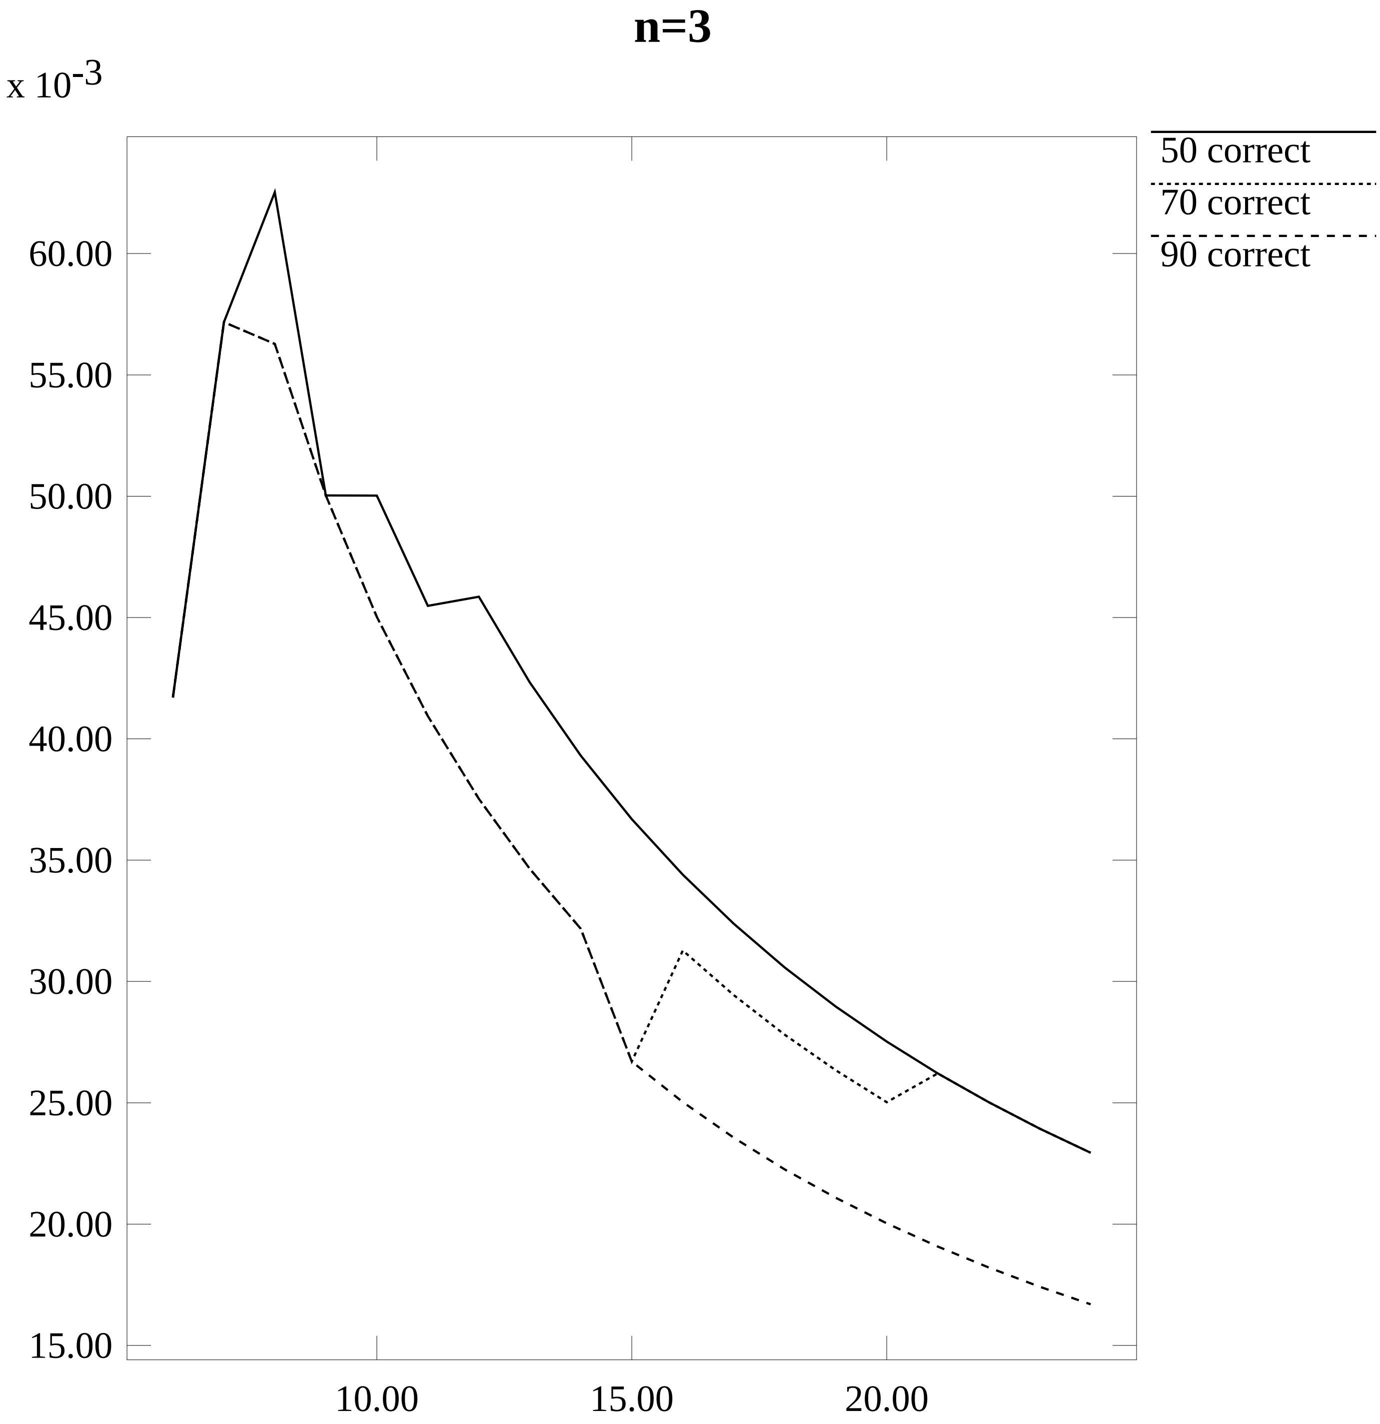
n=3
x 10^-3
60.00
55.00
50.00
45.00
40.00
35.00
30.00
25.00
20.00
15.00
10.00
15.00
20.00
50 correct
70 correct
90 correct

n=3

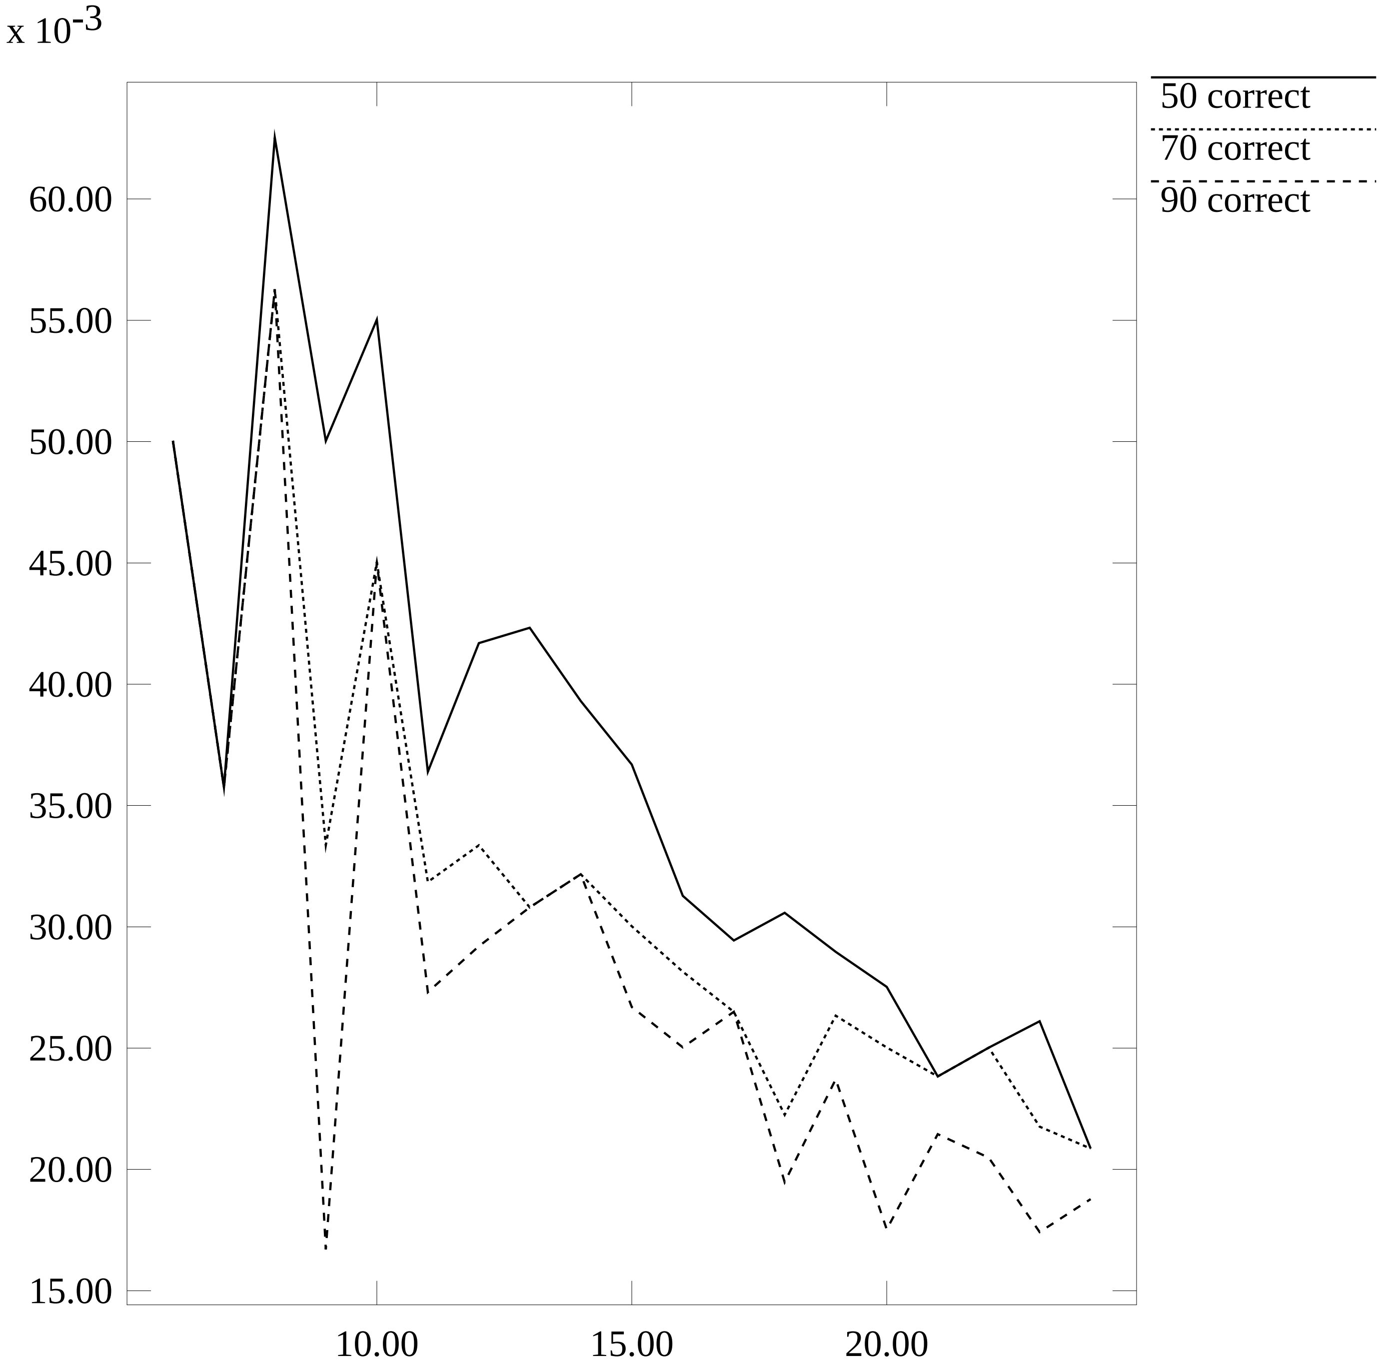
x 10^-3
60.00
55.00
50.00
45.00
40.00
35.00
30.00
25.00
20.00
15.00
10.00
15.00
20.00
50 correct
70 correct
90 correct

nc=1

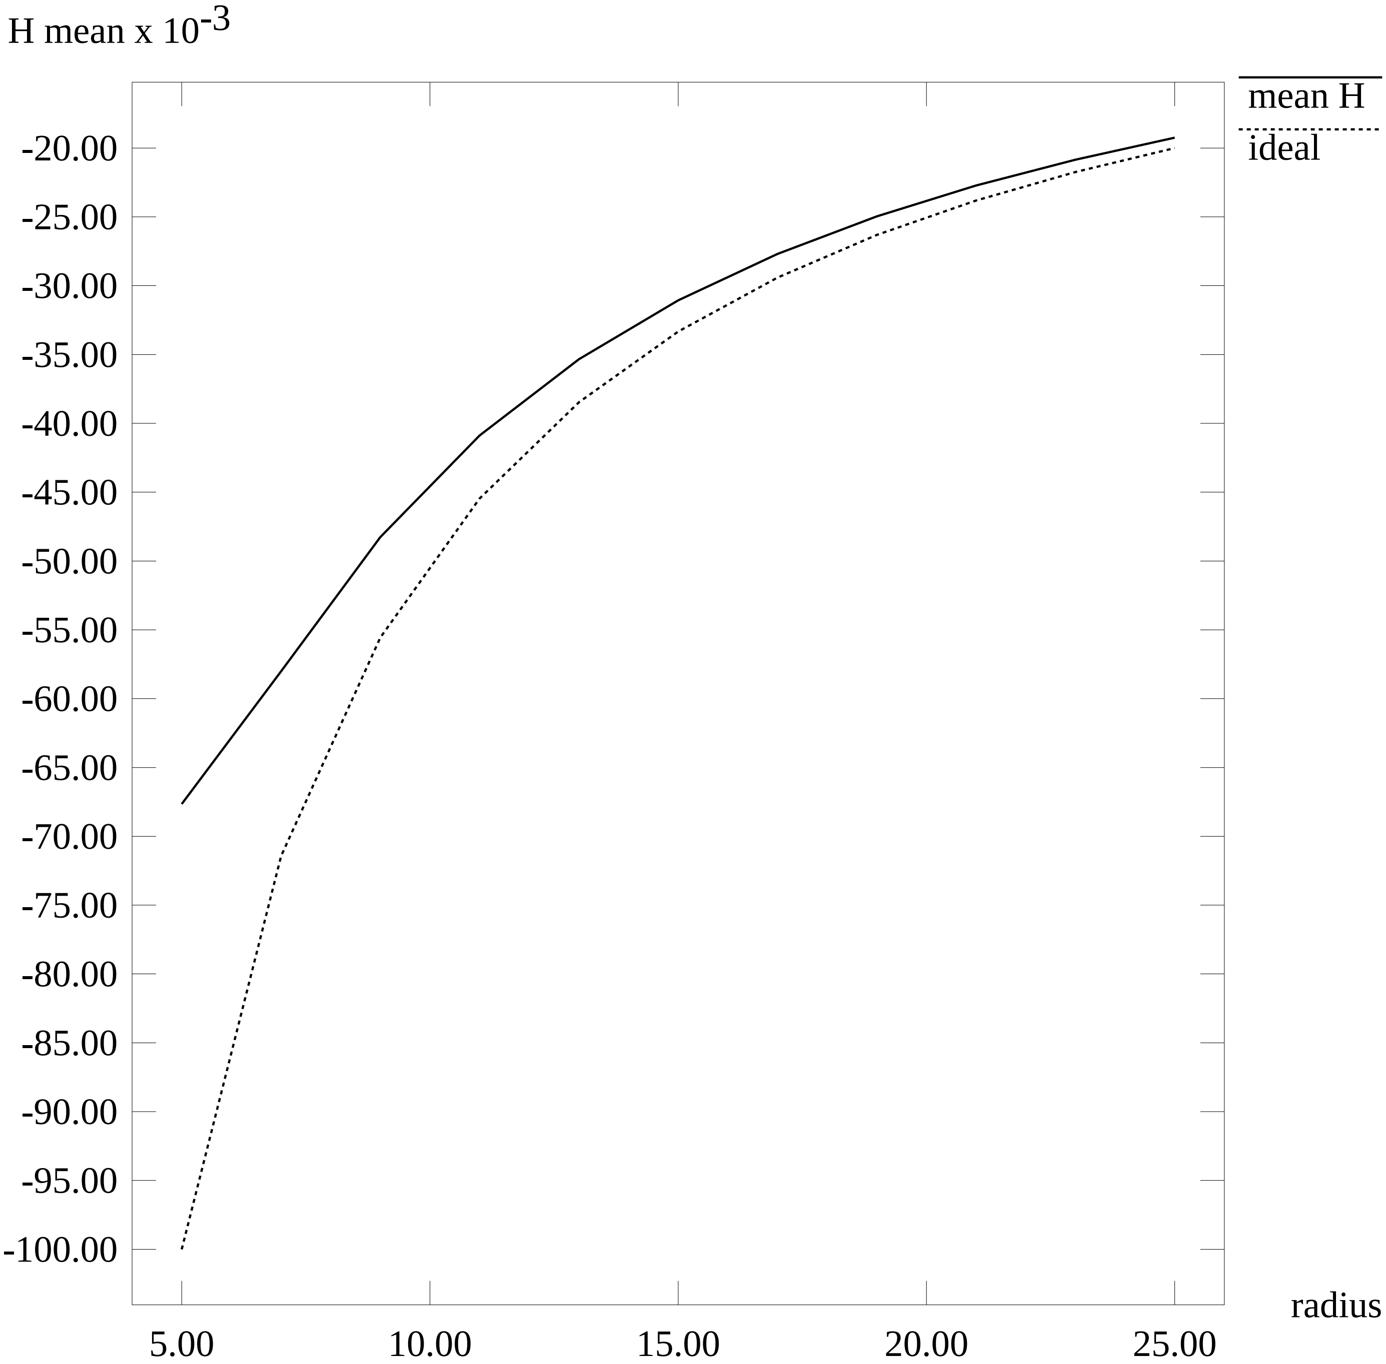
H mean x 10^{-3}
mean H
ideal
-20.00
-25.00
-30.00
-35.00
-40.00
-45.00
-50.00
-55.00
-60.00
-65.00
-70.00
-75.00
-80.00
-85.00
-90.00
-95.00
-100.00
radius
5.00
10.00
15.00
20.00
25.00

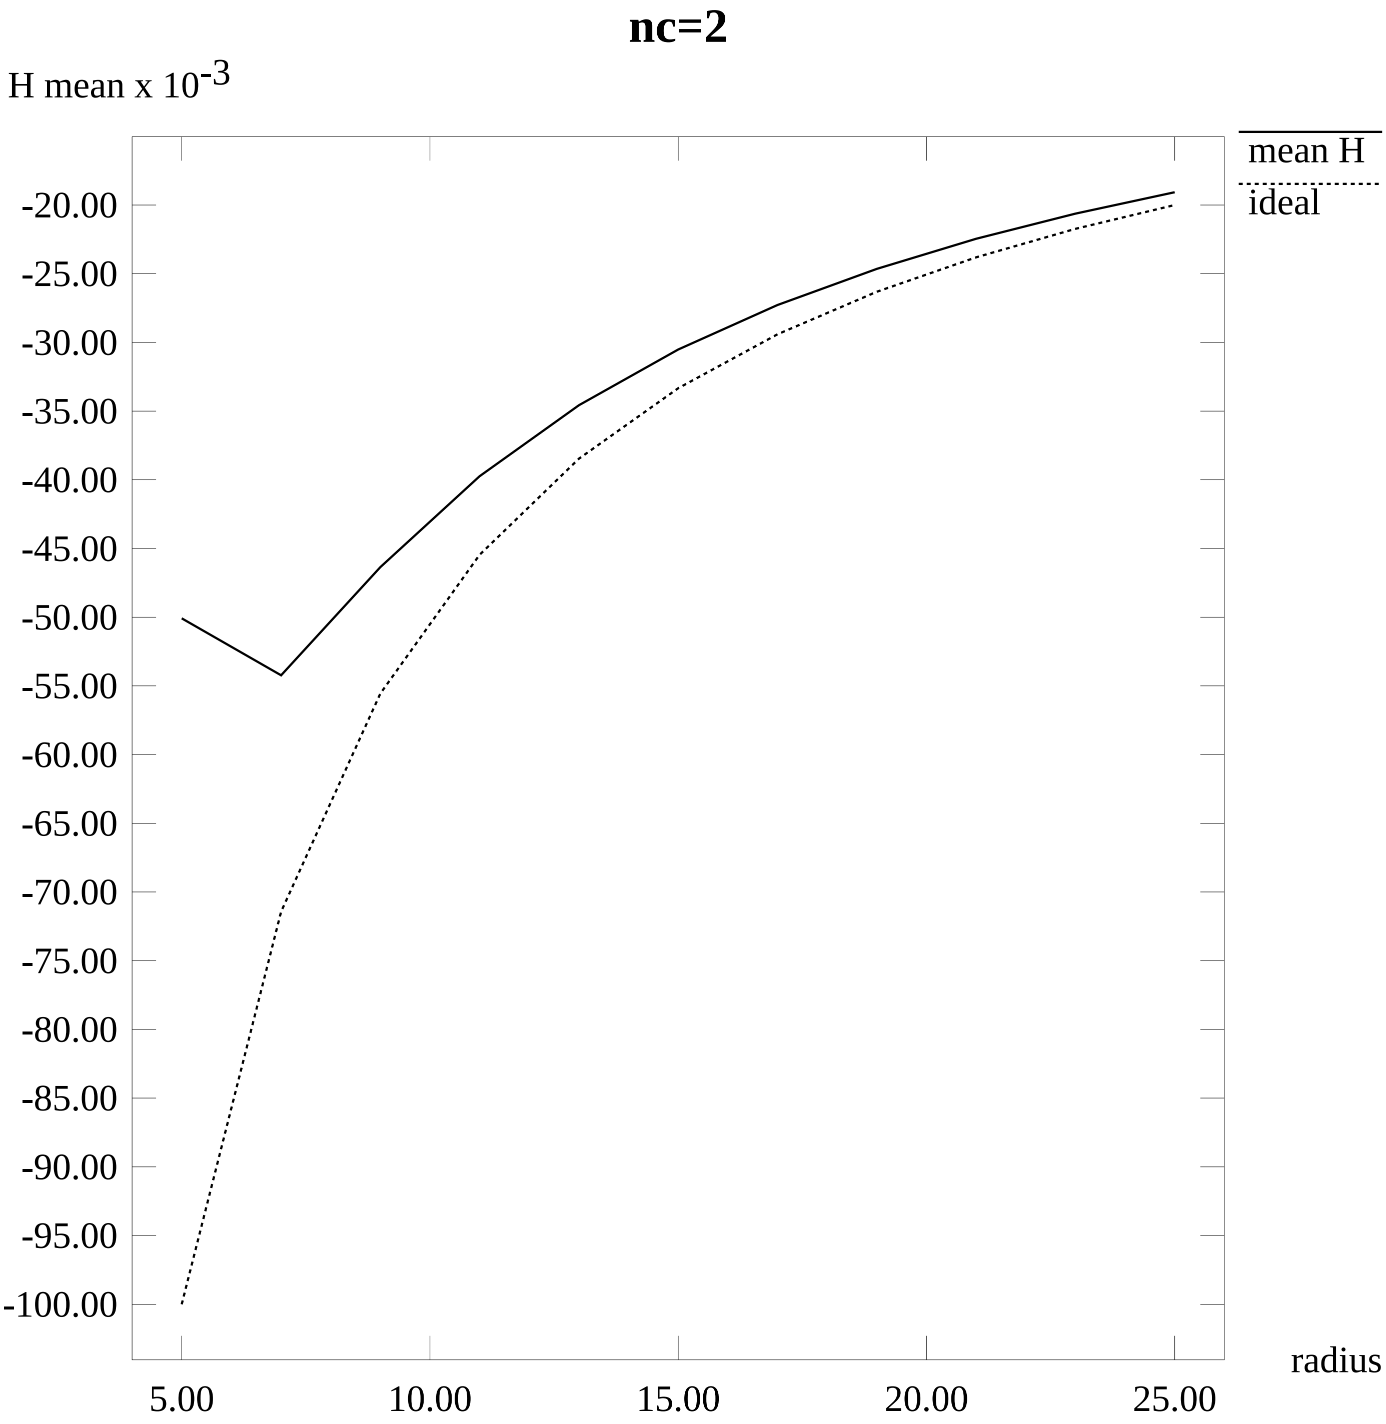
nc=2
H mean x 10^{-3}
-20.00
-25.00
-30.00
-35.00
-40.00
-45.00
-50.00
-55.00
-60.00
-65.00
-70.00
-75.00
-80.00
-85.00
-90.00
-95.00
-100.00
5.00
10.00
15.00
20.00
25.00
radius
mean H
ideal

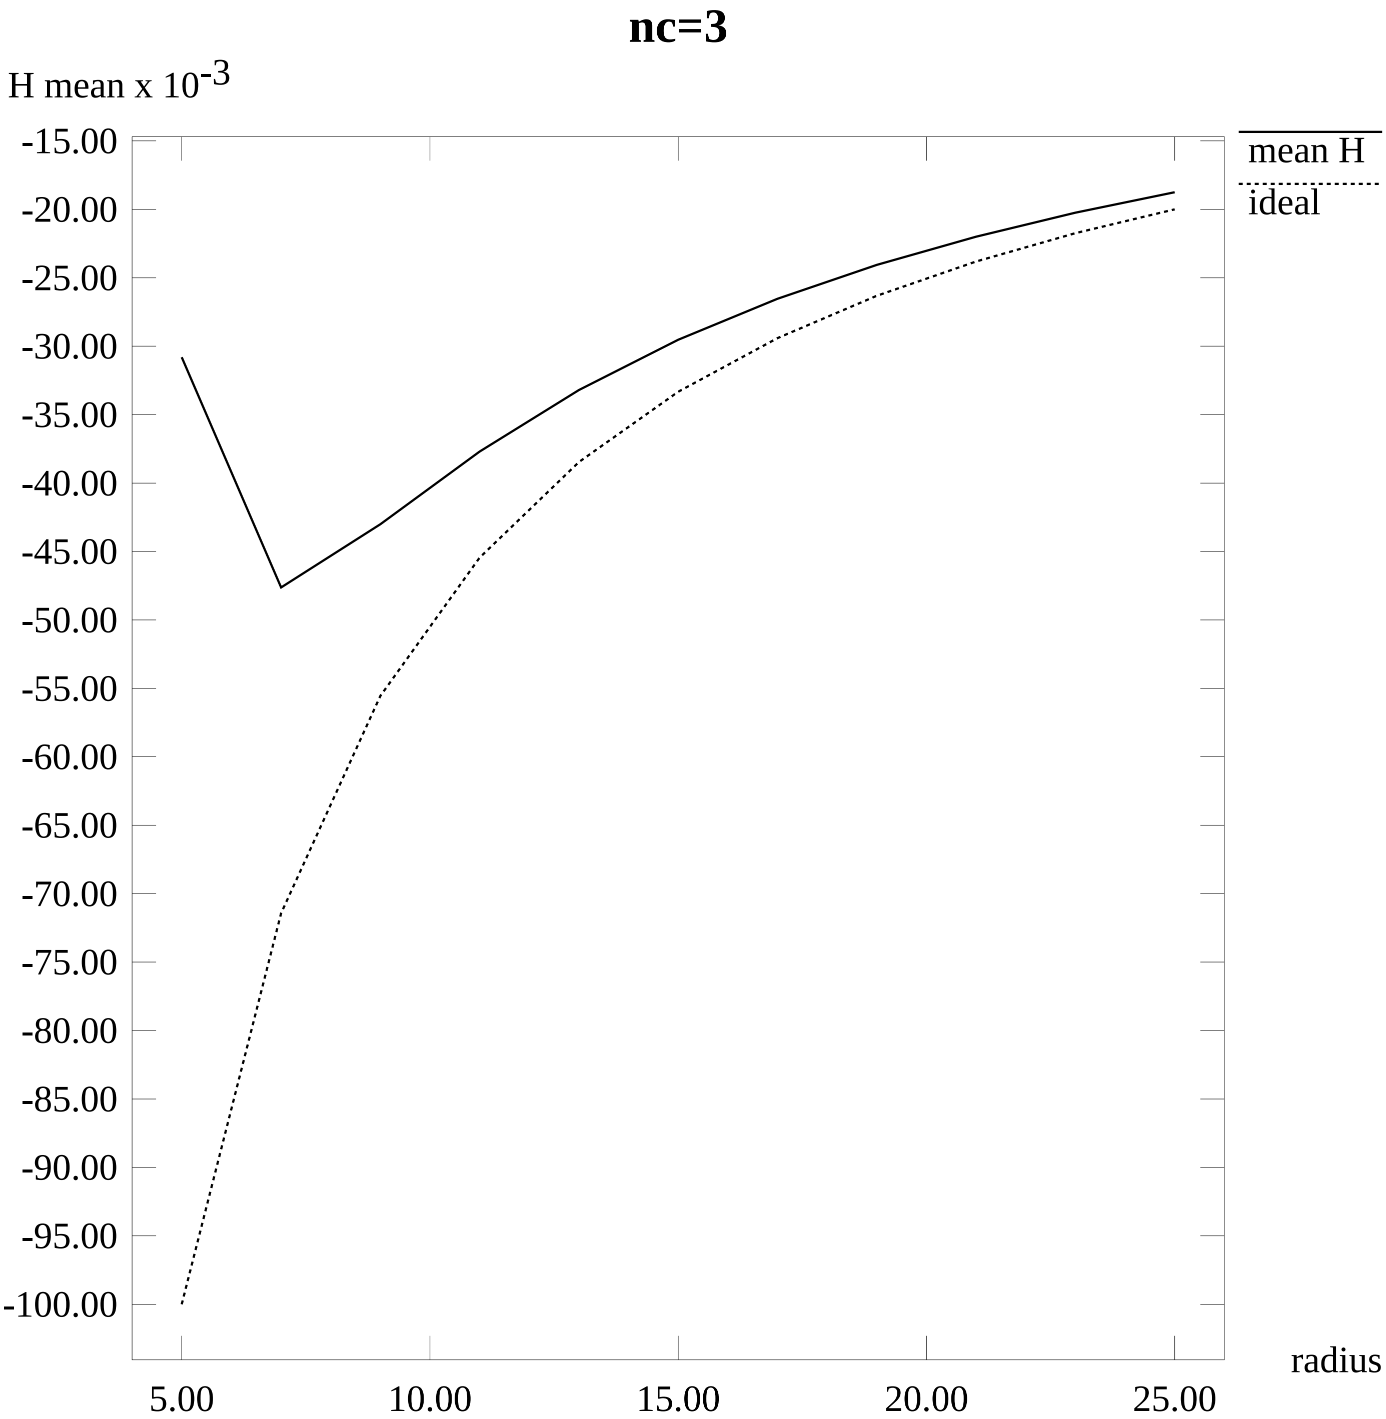
nc=3
H mean x 10-3
-15.00
-20.00
-25.00
-30.00
-35.00
-40.00
-45.00
-50.00
-55.00
-60.00
-65.00
-70.00
-75.00
-80.00
-85.00
-90.00
-95.00
-100.00
mean H
ideal
radius
5.00
10.00
15.00
20.00
25.00

nc=1

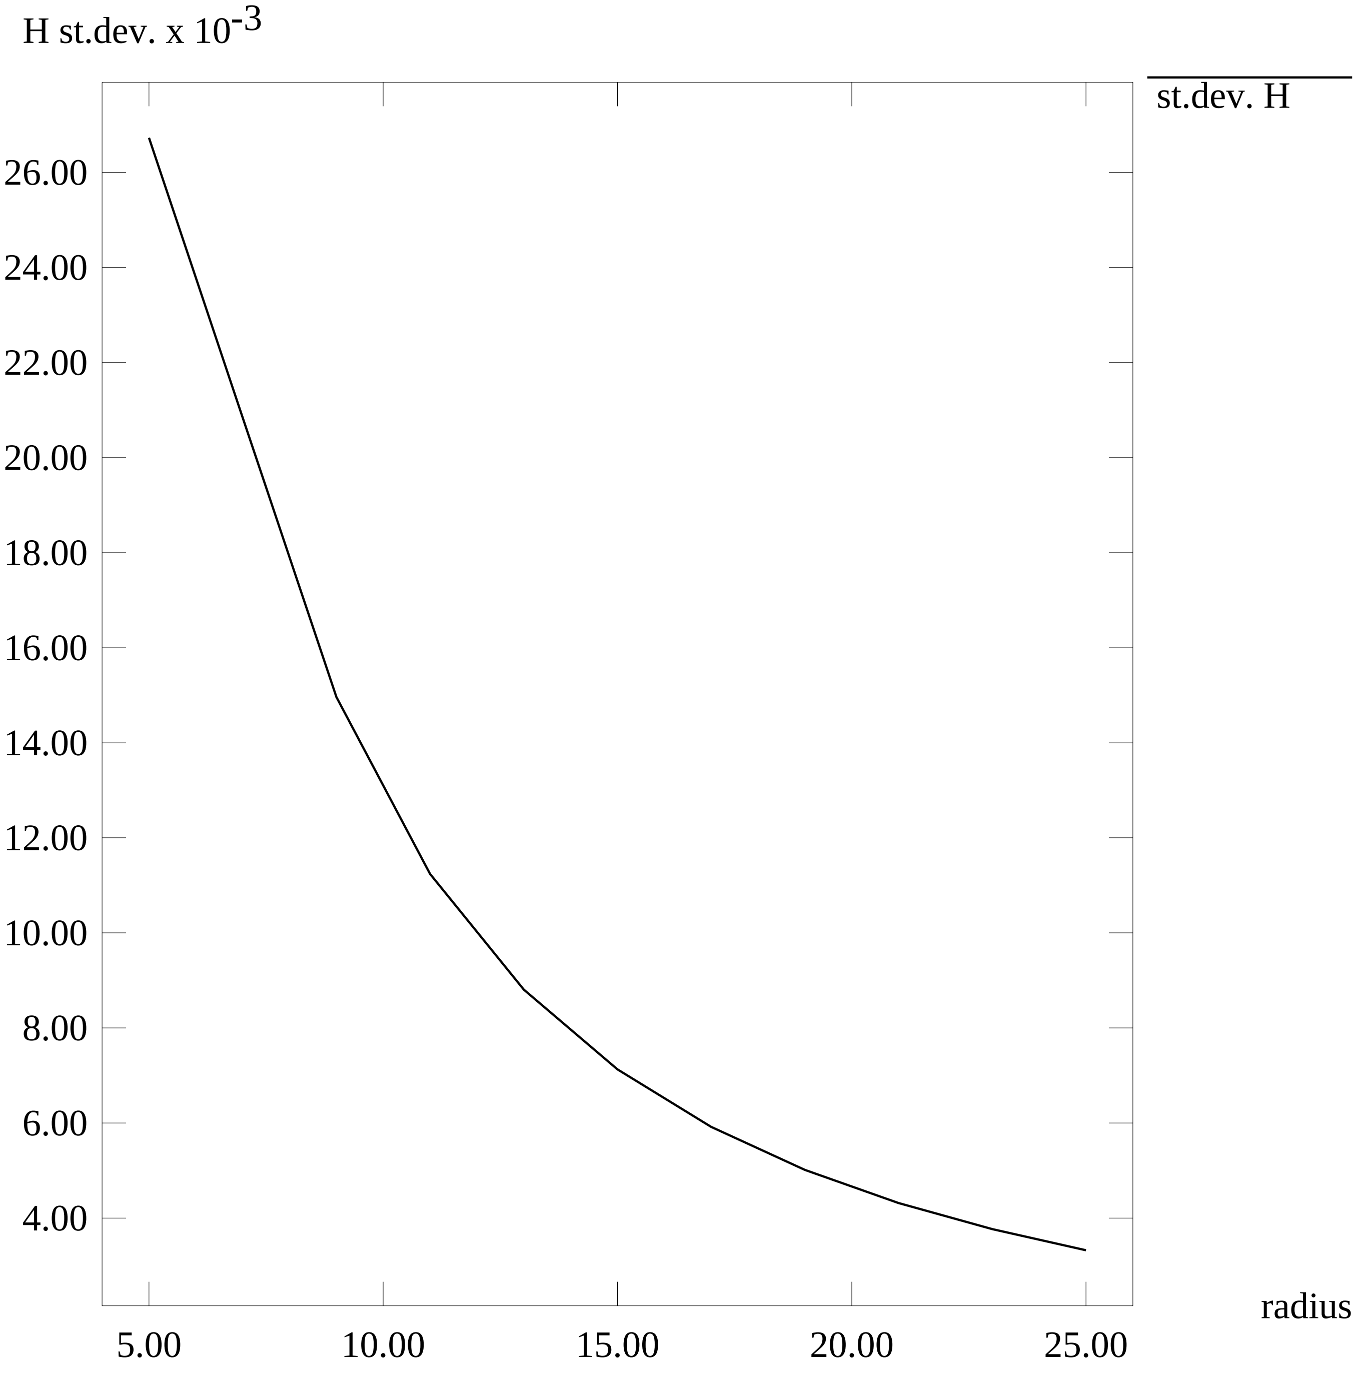

H st.dev. x 10^{-3}
st.dev. H
26.00
24.00
22.00
20.00
18.00
16.00
14.00
12.00
10.00
8.00
6.00
4.00
5.00
10.00
15.00
20.00
25.00
radius

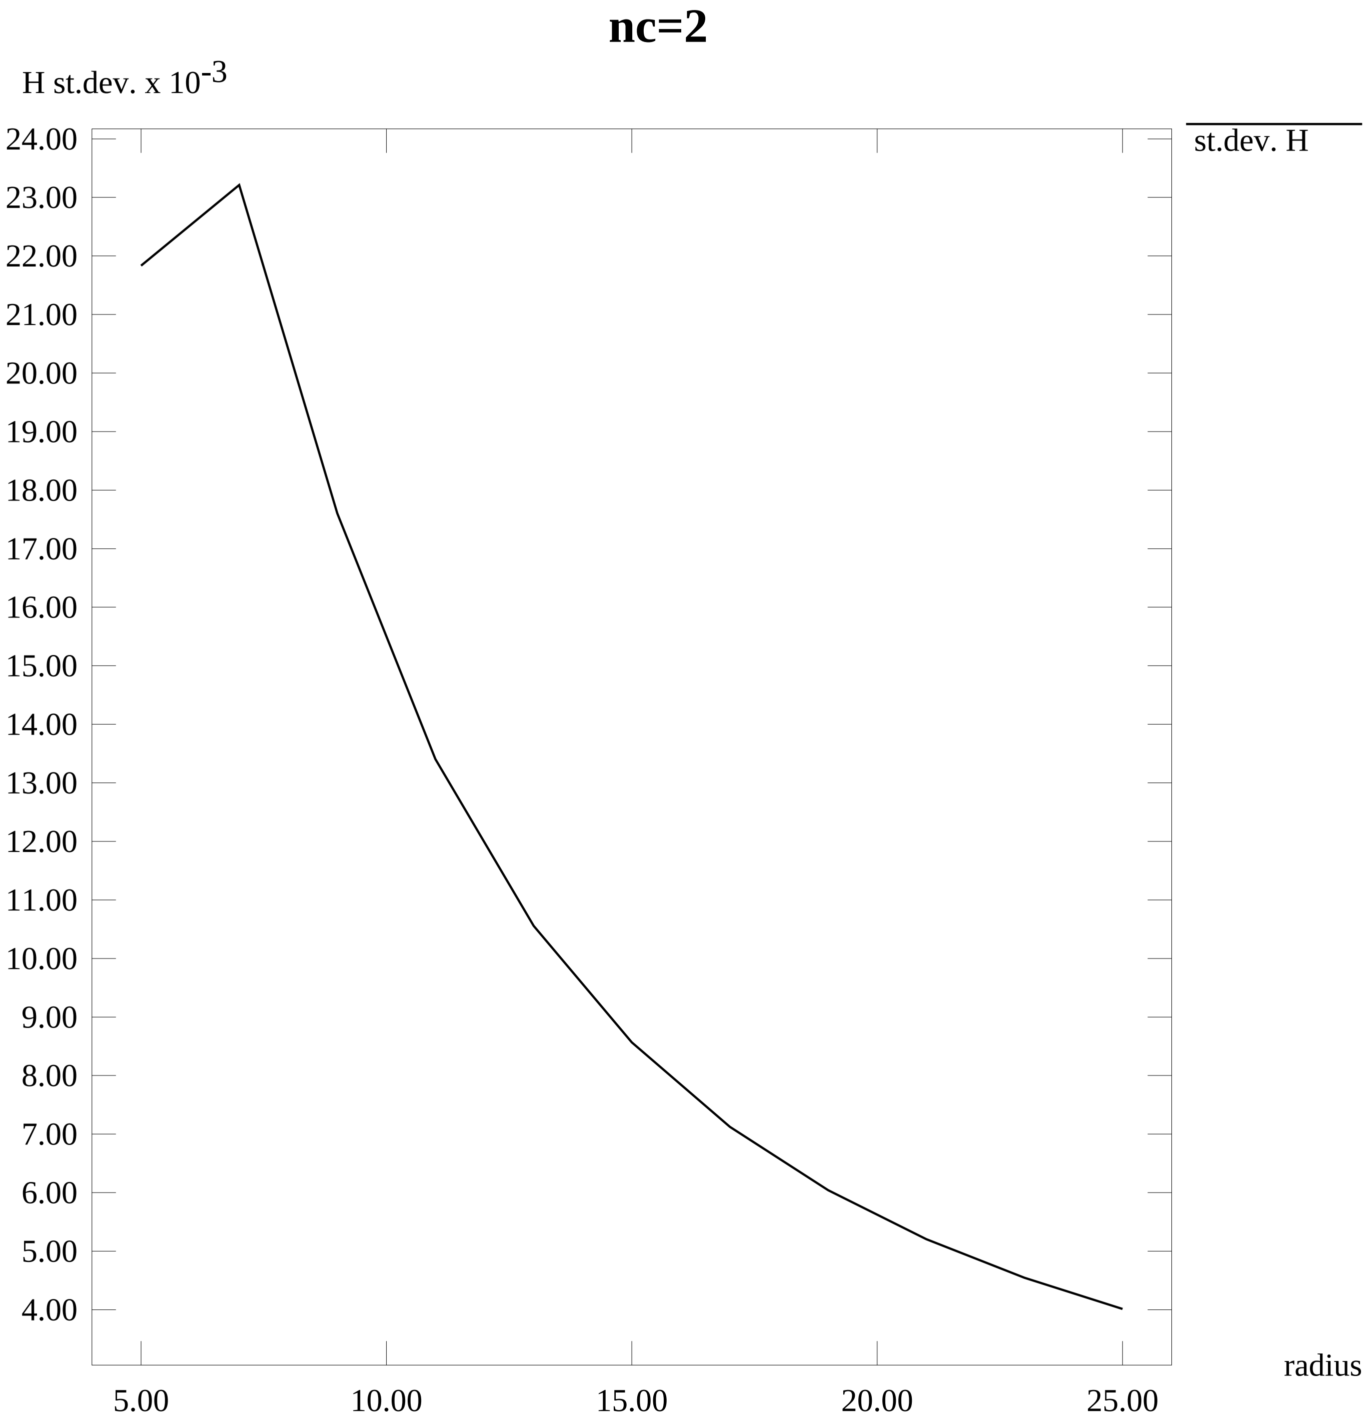
nc=2
H st.dev. x 10^{-3}
st.dev. H
24.00
23.00
22.00
21.00
20.00
19.00
18.00
17.00
16.00
15.00
14.00
13.00
12.00
11.00
10.00
9.00
8.00
7.00
6.00
5.00
4.00
radius
5.00
10.00
15.00
20.00
25.00

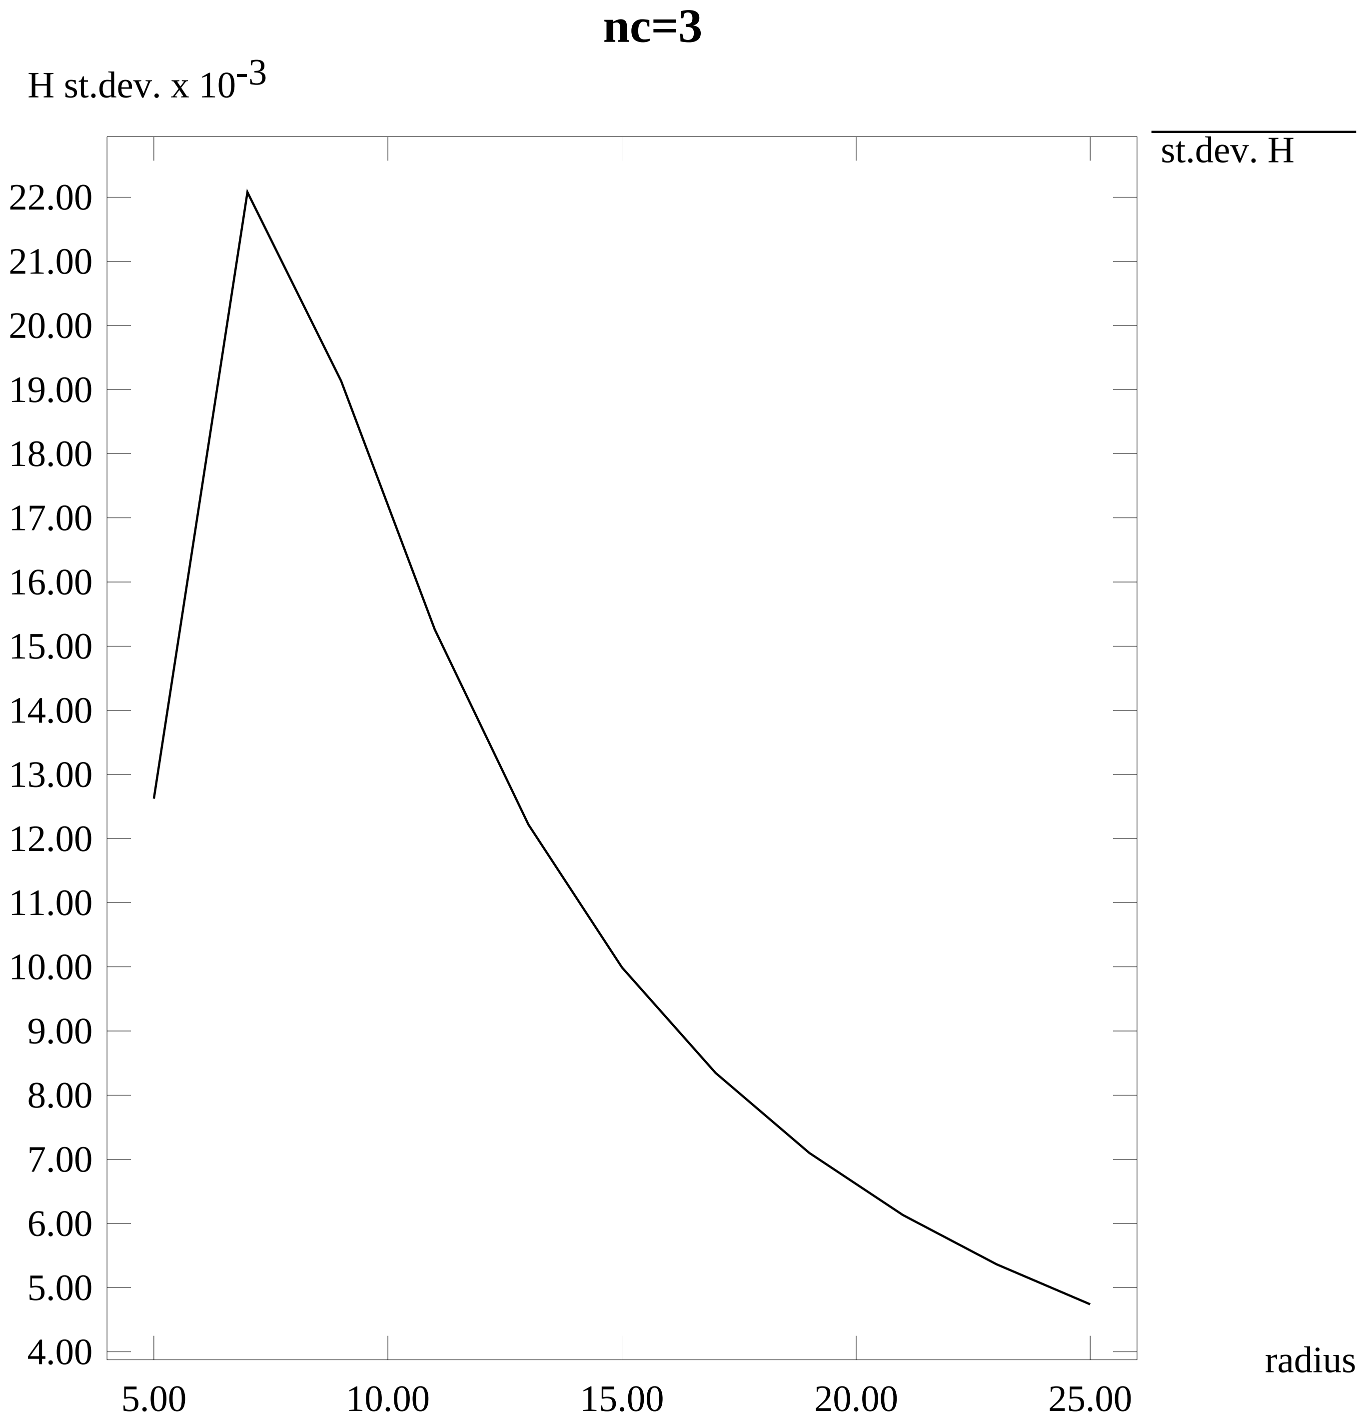

nc=3
H st.dev. x 10-3
st.dev. H
22.00
21.00
20.00
19.00
18.00
17.00
16.00
15.00
14.00
13.00
12.00
11.00
10.00
9.00
8.00
7.00
6.00
5.00
4.00
5.00
10.00
15.00
20.00
25.00
radius