Nonlinear adaptively learned optimization for object localization in 3D medical images

Mayalen Etcheverry, Bogdan Georgescu, Benjamin Odry, Thomas J. Re, Shivam Kaushik, Bernhard Geiger, Nadar Mariappan, Sasa Grbic, and Dorin Comaniciu

Siemens Healthineers, Digital Services, Digital Technology & Innovation

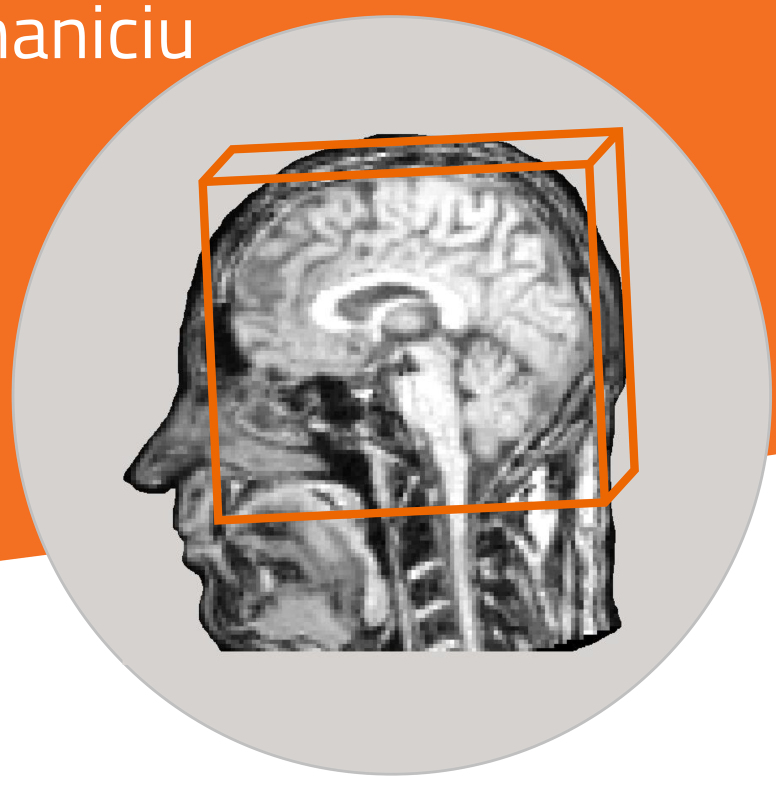

1 INTRODUCTION

Automatic object localization is an important prerequisite for many tasks in medical imaging analysis such as image registration, organ segmentation, lesion quantification and abnormality detection. To this end, we propose to automatically localize a 3D object by estimating the 9 pose parameters for position, orientation and size.

We introduce a deep reinforcement learning-based method, where an intelligent agent actively learns to optimize the pose parameters by performing a sequence of simple control actions. An adaptive sequential search across different scale representations of the environment is proposed.

We build upon the work of Ghesu et al. [1] that uses reinforcement learning to identify the location of an anatomical landmark. We propose to extend the method to a wider range of image analysis applications by expanding the search space to an nonlinear multi-dimensional parametric space.

2 METHOD

1- Object Localization as a Markov Decision Process (S, A, p, R, γ):

- State representation s: content of the current region occupied by the agent (*visible region*) + a fixed margin of voxels (*additional context*)
- Control actions a: 2D move actions (modify the current object geometry) + one stop action (terminate the search)
- Reward function r: $r_t = \begin{cases} dist(x_t, x^*) - dist(x_{t+1}, x^*) \ if \ a_t \in \{1, \dots, 2D\} \\ \left(\frac{dist(x_t, x^*) - d_{min}}{d_{max} - d_{min}} - 0.5\right) * 6 \ if \ a_t = 2D + 1 \\ -1 \ \ if \ \mathrm{s}_{t+1} \text{non legal state} \end{cases}$
- Goal: maximize the cumulative future reward $R = \sum_{t=0}^{T} \gamma^t r^t$

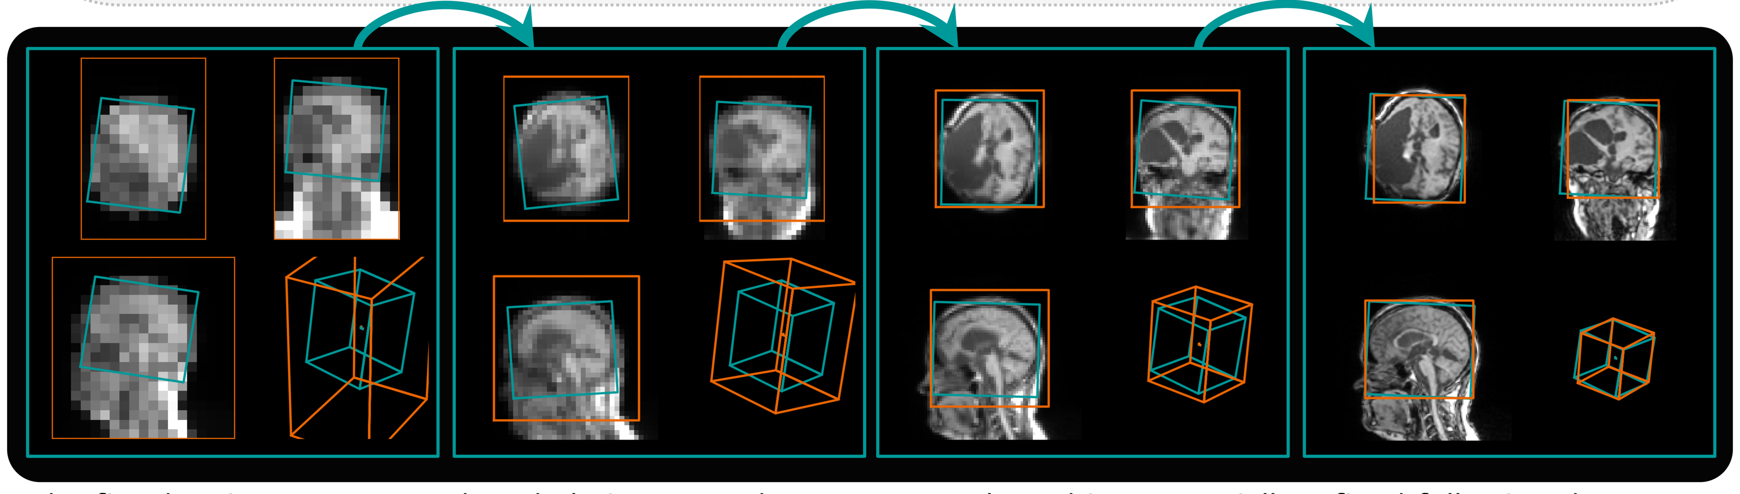

The first box is set to cover the whole image at the coarsest scale and is sequentially refined following the agent's decisions. The final agent position is taken as localization result.

2- Deep Q-learning to find the optimal policy π^*

- The optimal policy returns the action maximizing the cumulative discounted reward: $\forall \, s \in S : \pi^* = \underset{a}{\operatorname{argmax}} \, Q^*(s, a)$ where
  $Q^*(s, a) = \max_{\pi} \mathbb{E}[\sum_{t=0}^{T} \gamma^t r^t \, | s_0 = s, a_0 = a, \pi]$
- Bellman equation: $Q^*(s, a) = \mathbb{E}_{s' \sim \varepsilon}[r + \gamma \max_{a'} Q^*(s', a') \; | s, a]$ (1)
- DQN[2]: use of a NN function approximation $Q^*(s, a) \approx Q(s, a, \theta)$
  The training minimizes a sequence of loss functions
  $\mathrm{L_i}(\theta_i) = \mathbb{E}_{s,a,r,s'}(y_i - \mathrm{Q}(\mathrm{s}, \mathrm{a}; \theta_i))^2$ where
  $y_i = r + \gamma Q_{target}(s', \underset{a'}{\operatorname{argmax}} \, Q(s', a'; \theta_i); \theta^-)$
  expressing how far $\mathrm{Q}(\mathrm{s}, \mathrm{a}; \theta_i)$ is from its target y_i (1)
- Experience replay, guided ε-greedy exploration, double Q-learning [3]

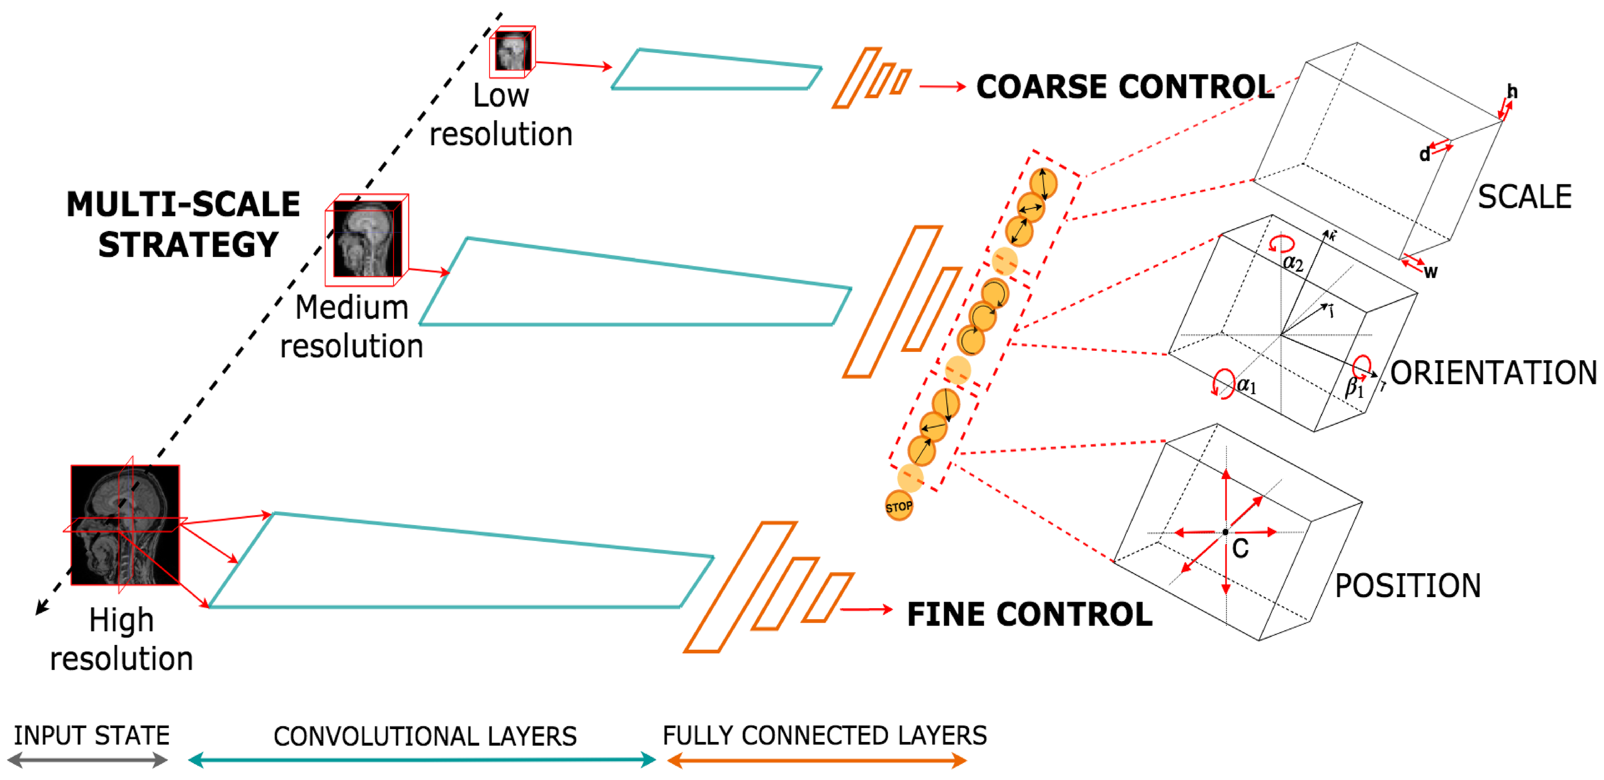

3- Multi-Scale Progressive Control Strategy

Progressive spanning-scheme of the nonlinear D-dimensional search space:

- Discretization of the context (continuous volumetric image) into a multi-scale image pyramid with increasing image resolution
- Discretization of the D-dimensional parametric search space with progressive incremental precision across scales

3 RESULTS

Experiment purpose: Standardize orientations of acquisitions by localizing a standard box from Scout / Localizer MRI scans of the head region.

Dataset: 500 training + 30 testing + 15 challenging cases (tumors or fluid swelling, in plane rotation of the head, cropped top of the skull). Annotated by experts following a specific procedure.

Evaluation:

- Angular measures $(\alpha_1, \alpha_2, \beta_1)$ and distance measures (for each face $\delta_R, \delta_L, \delta_A, \delta_P, \delta_I, \delta_S$) between detected and ground truth boxes.

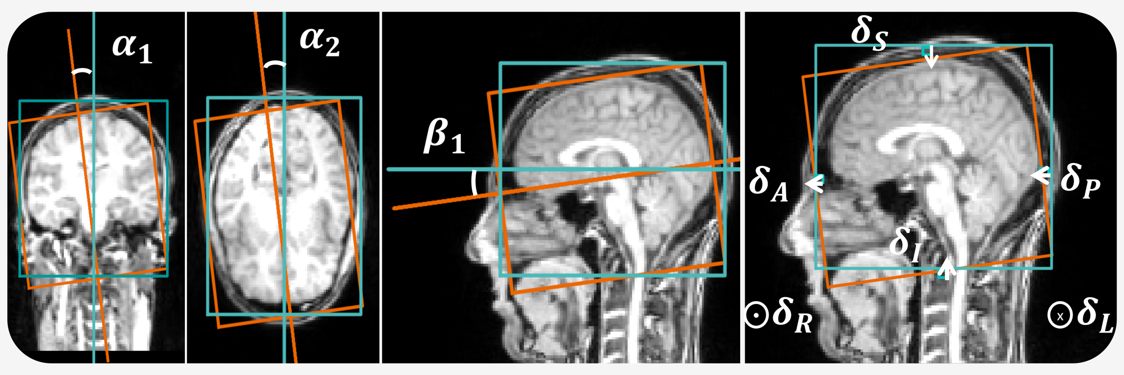

- Comparison with inter-rater variability and a previous landmark-based method.

Detection error:

	Inter-rater	Landmark-based	Our approach (4mm)	Our approach (2mm)
$\alpha_1(°)$	0.99(≤3.50)	0.92(≤3.45)	1.28(≤3.78)	**0.92(≤3.23)**
$\alpha_2(°)$	1.04(≤4.71)	0.99(≤4.93)	1.20(≤4.46)	**0.97(≤2.11)**
$\beta_1(°)$	1.47(≤5.19)	2.00(≤6.86)	1.62(≤6.35)	**1.39(≤5.86)**
δ_R(mm)	1.32(≤3.54)	2.06(≤5.78)	2.65(≤7.54)	**1.45(≤3.30)**
δ_L(mm)	1.45(≤4.75)	1.89(≤5.03)	2.20(≤8.68)	**1.83(≤4.95)**
δ_A(mm)	2.00(≤3.36)	**1.65(≤4.93)**	2.46(≤6.07)	1.94(≤6.08)
δ_P(mm)	1.48(≤3.89)	1.86(≤9.62)	3.31(≤9.68)	**1.65(≤5.68)**
δ_I(mm)	3.33(≤3.61)	**2.22(≤6.00)**	3.12(≤11.5)	2.74(≤8.21)
δ_S(mm)	1.3(≤3.28)	**2.13(≤5.74)**	3.04(≤7.46)	2.16(≤6.31)

Robustness: No major failure cases on the challenging cases

Runtime: 0.6 secs on GPU (< 0.15 secs if stopped at 4mm scale-level)

4 CONCLUSION

- Novel approach to sequentially search for a target object inside 3D medical images
- Achieves **high speed** and **high accuracy** results.
- The methodology can learn optimization strategies **eliminating the need for exhaustive search or for complex generic nonlinear optimization techniques.**
- Method can be applied to a broad range of problems.

References:

[1] Ghesu, F.C., Georgescu, B., Zheng, Y., Grbic, S., Maier, A., Hornegger, J. and Comaniciu, D.: Multi-Scale Deep Reinforcement Learning for Real-Time 3D-Landmark Detection in CT Scans. IEEE Transactions on Pattern Analysis and Machine Intelligence. (2017)

[2] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. and Riedmiller, M.: Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602. (2013)

[3] Van Hasselt, H., Guez, A. and Silver, D.: Deep Reinforcement Learning with Double Q-Learning. In AAAI, Vol. 16, pp. 2094{2100. (2016)

usa.siemens.com/abouthealthineers

SIEMENS Healthineers