Software Self-configurability in GPU-accelerated Robot Vision

Luigi Nardi, PhD
Imperial College London
Software Performance Optimisation group

@GTC-EU Amsterdam
September 29th 2016

In collaboration with:
B. Bodin, M Z. Zia, H. Wagstaff, D. Carroll, A. White, E. Vespa, S. Saeedi, G. S. Shenoy, M. K. Emani, J. Mawer, A. Nisbet, M. Luján, B. Franke, M. F. P. O'Boyle, A. J. Davison, P. H. J. Kelly and S. Furber

Luigi Nardi - Imperial College London

The three R's of vision: Spectrum of Computer Vision Research

Reconstruction

Scalable Kinect Fusion (2013)

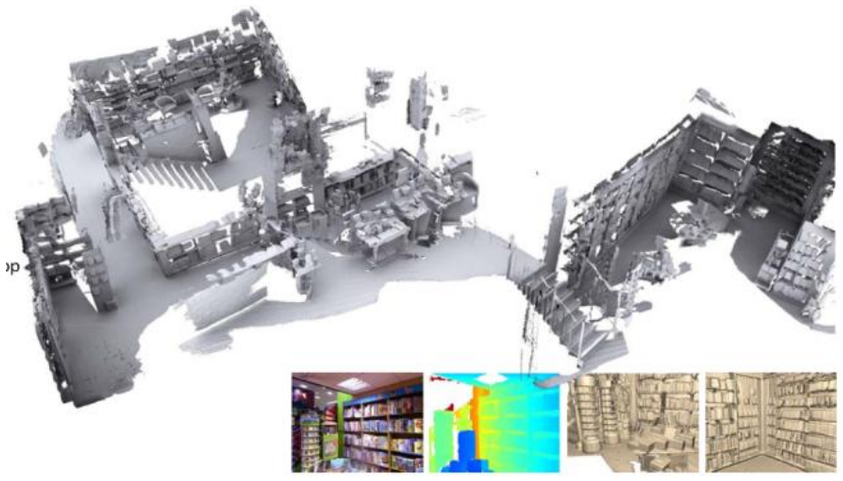

Building Rome on a cloudless day (2010)

Recognition

Deep learning for scalable Object class detection (2014)

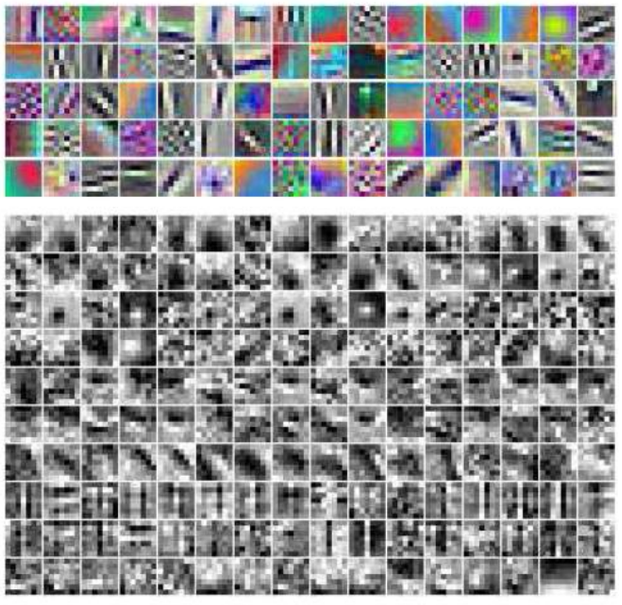

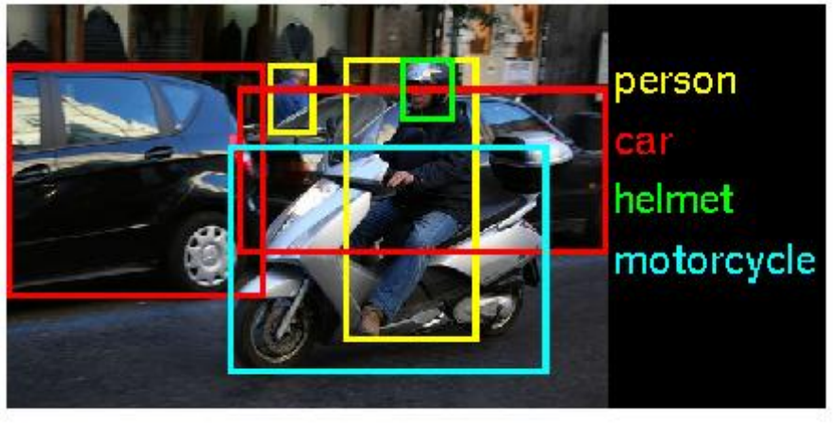

Reorganisation or Grouping

Contour detection and segmentation (2011)

Credit: Zeeshan Zia

Luigi Nardi - Imperial College London

Simultaneous localisation and mapping (SLAM)

Build a coherent world representation and localise the camera in real-time

Sparse SLAM

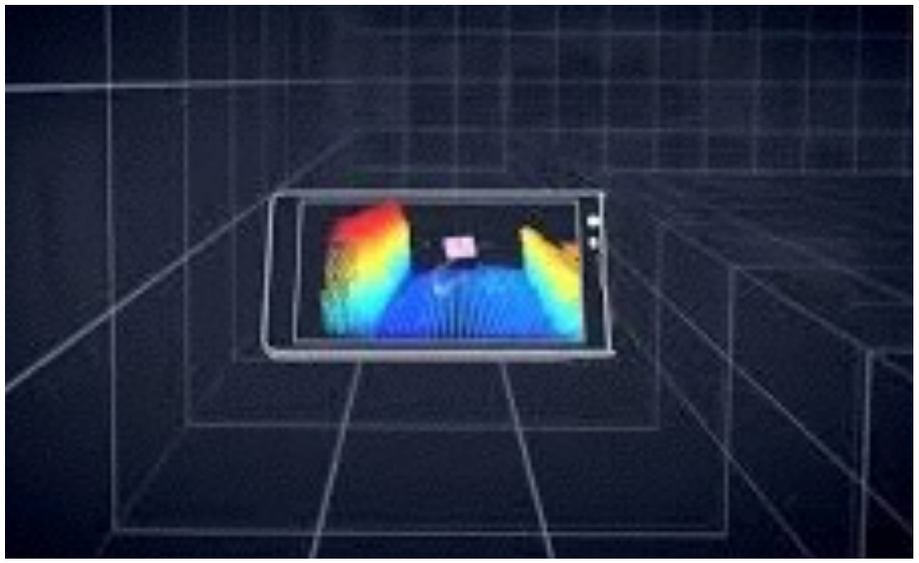

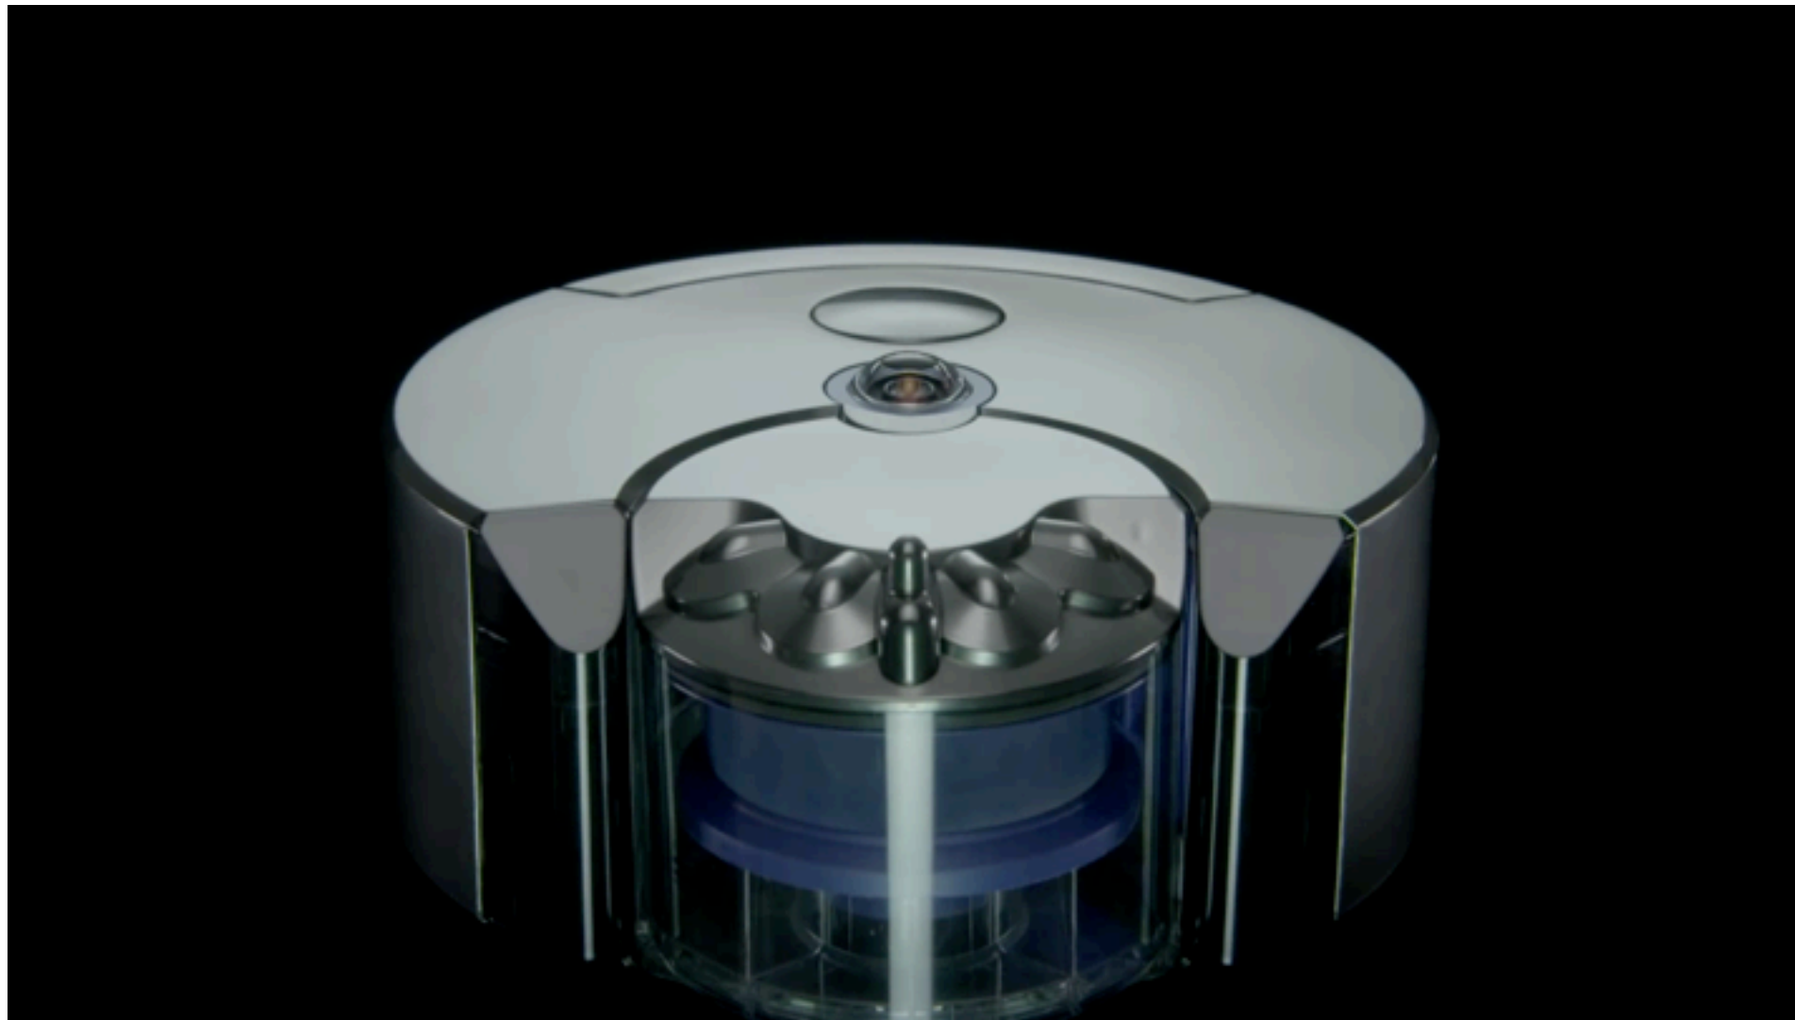

Video:
Dyson 360 Eye

SIGGRAPH Talks 2011

KinectFusion:

Real-Time Dynamic 3D Surface Reconstruction and Interaction

Shahram Izadi 1, Richard Newcombe 2, David Kim 1,3, Otmar Hilliges 1,
David Molyneaux 1,4, Pushmeet Kohli 1, Jamie Shotton 1,
Steve Hodges 1, Dustin Freeman 5, Andrew Davison 2, Andrew Fitzgibbon 1

1 Microsoft Research Cambridge 2 Imperial College London
3 Newcastle University 4 Lancaster University
5 University of Toronto

Dense SLAM

Video: **KinectFusion**
[Newcombe et al. ISMAR 2011]

Luigi Nardi - Imperial College London

Simultaneous localisation and mapping (SLAM)

Build a coherent world representation and localise the camera in real-time

Dense SLAM

In this talk I will focus on two dense algorithms:

- KinectFusion [Newcombe et al. ISMAR 2011]
- ElasticFusion [Whelan et al. RSS 2015]

Jesse Clayton (NVIDIA)
3D reconstruction

Applications, e.g.:

- Robotics
- Autonomous driving
- 3D printing
- Augmented reality
- Telepresence

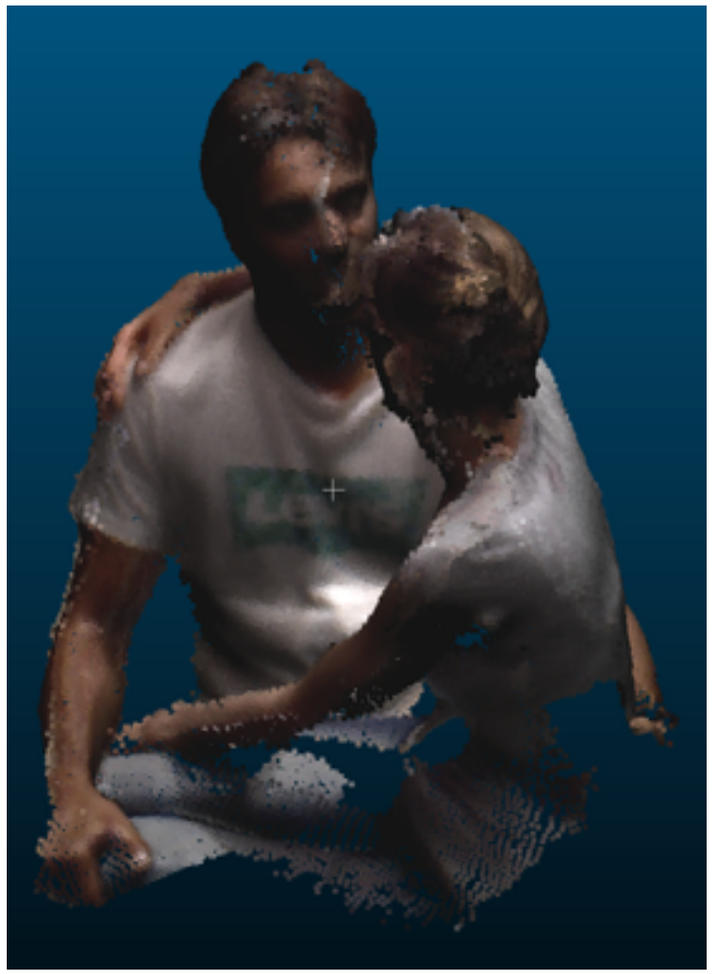

Daniele and Daniela
3D reconstruction

Pictures featured using ElasticFusion

What CV researchers say about **KinectFusion** and **ElasticFusion** performance

"Cannot run in real-time on mobile"

"You need a fat GPU to run it"

Holistic approach to SLAM performance:

SLAMBench

Computer Vision

Compiler/Runtime

Kernels → **Architectures** → **Correctness** → **Performance** → **Metrics**

```
bilateralFilter (..)
halfSampleRobust (..)
renderVolume (..)
integrate (..)
      :
      :
```

Hardware

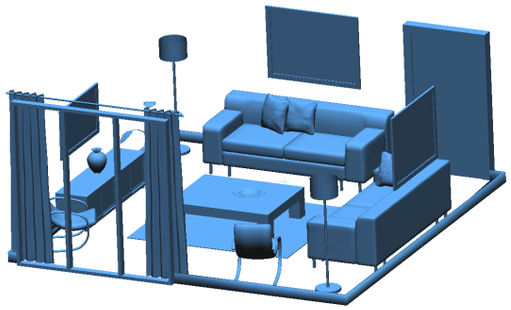

ICL-NUIM Dataset

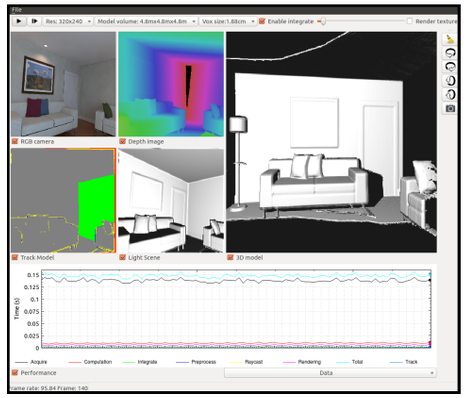

Frame rate

Energy

Accuracy

A publicly-available benchmarking framework for quantitative, comparable and validatable experimental research to investigate trade-offs in performance, accuracy and energy consumption of a SLAM system

Error metric: absolute trajectory error (ATE) based on dataset ground truth

Introducing SLAMBench, a performance and accuracy benchmarking methodology for SLAM **(ICRA 2015)**

SLAMBench framework

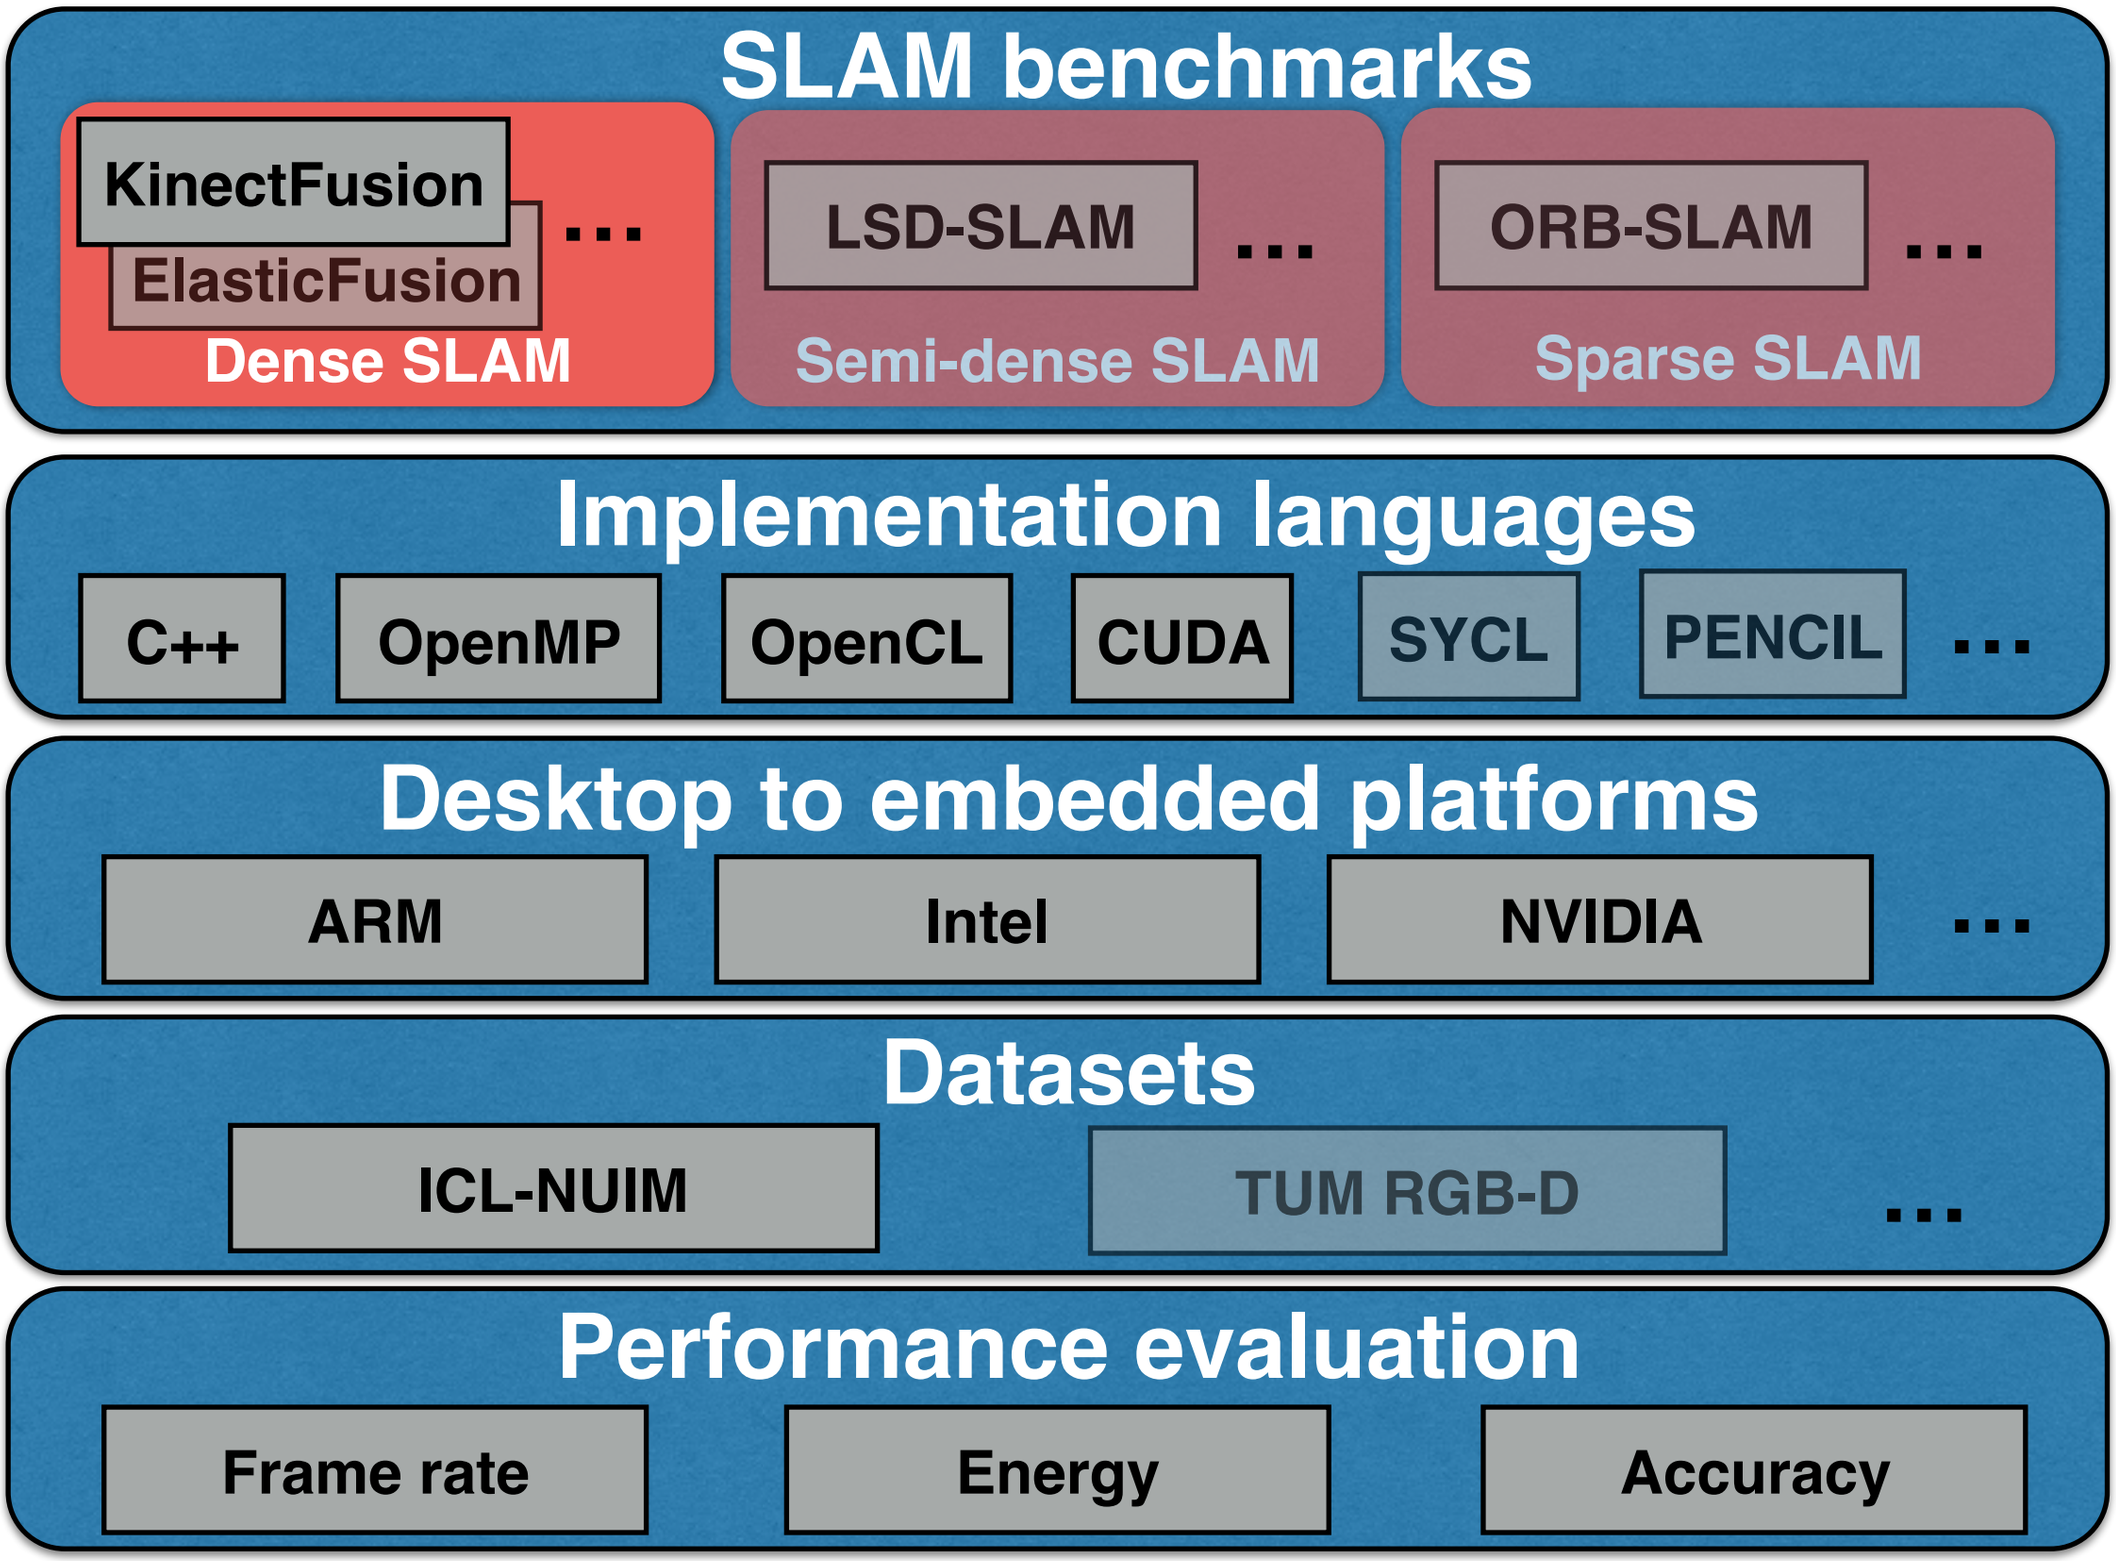

KinectFusion algorithmic features

Features	Ranges
Volume resolution	64x64x64, 128x128x128, 256x256x256, 512x512x512
μ distance	0 .. 0.5
Pyramid level iterations (3 levels)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11
Image resolution (image ratio)	1, 2, 4, 8
Tracking rate	1, 2, 3, 4, 5
ICP threshold	10^{-6} .. 10^{2}
Integration rate	1 .. 30

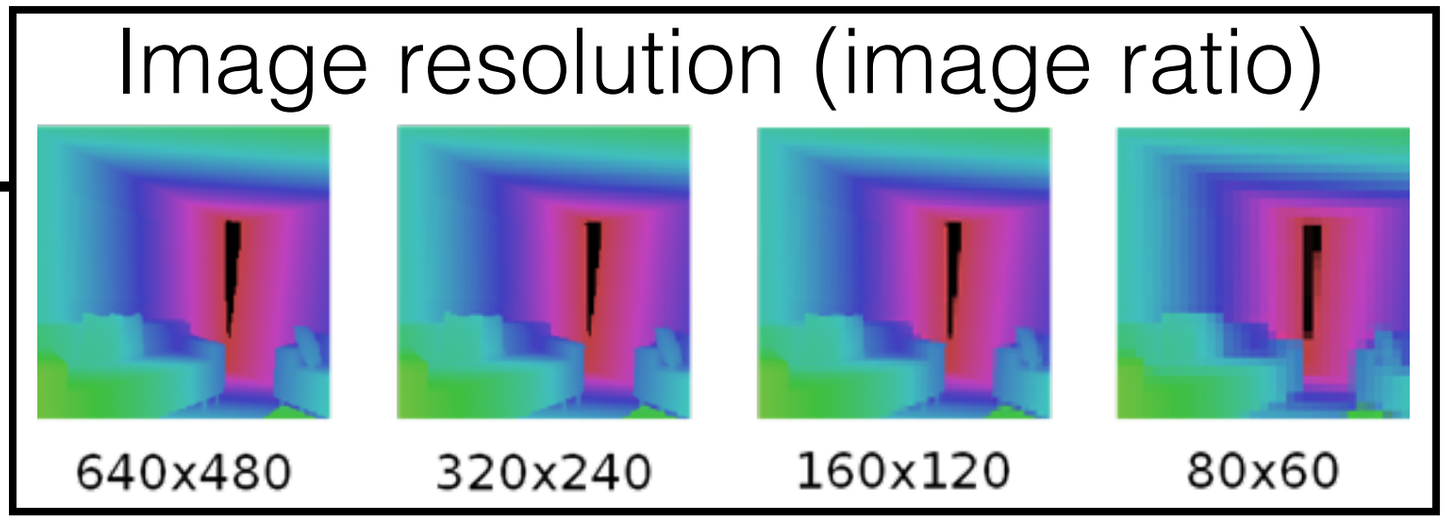

Different algorithmic features for ElasticFusion

Motivation

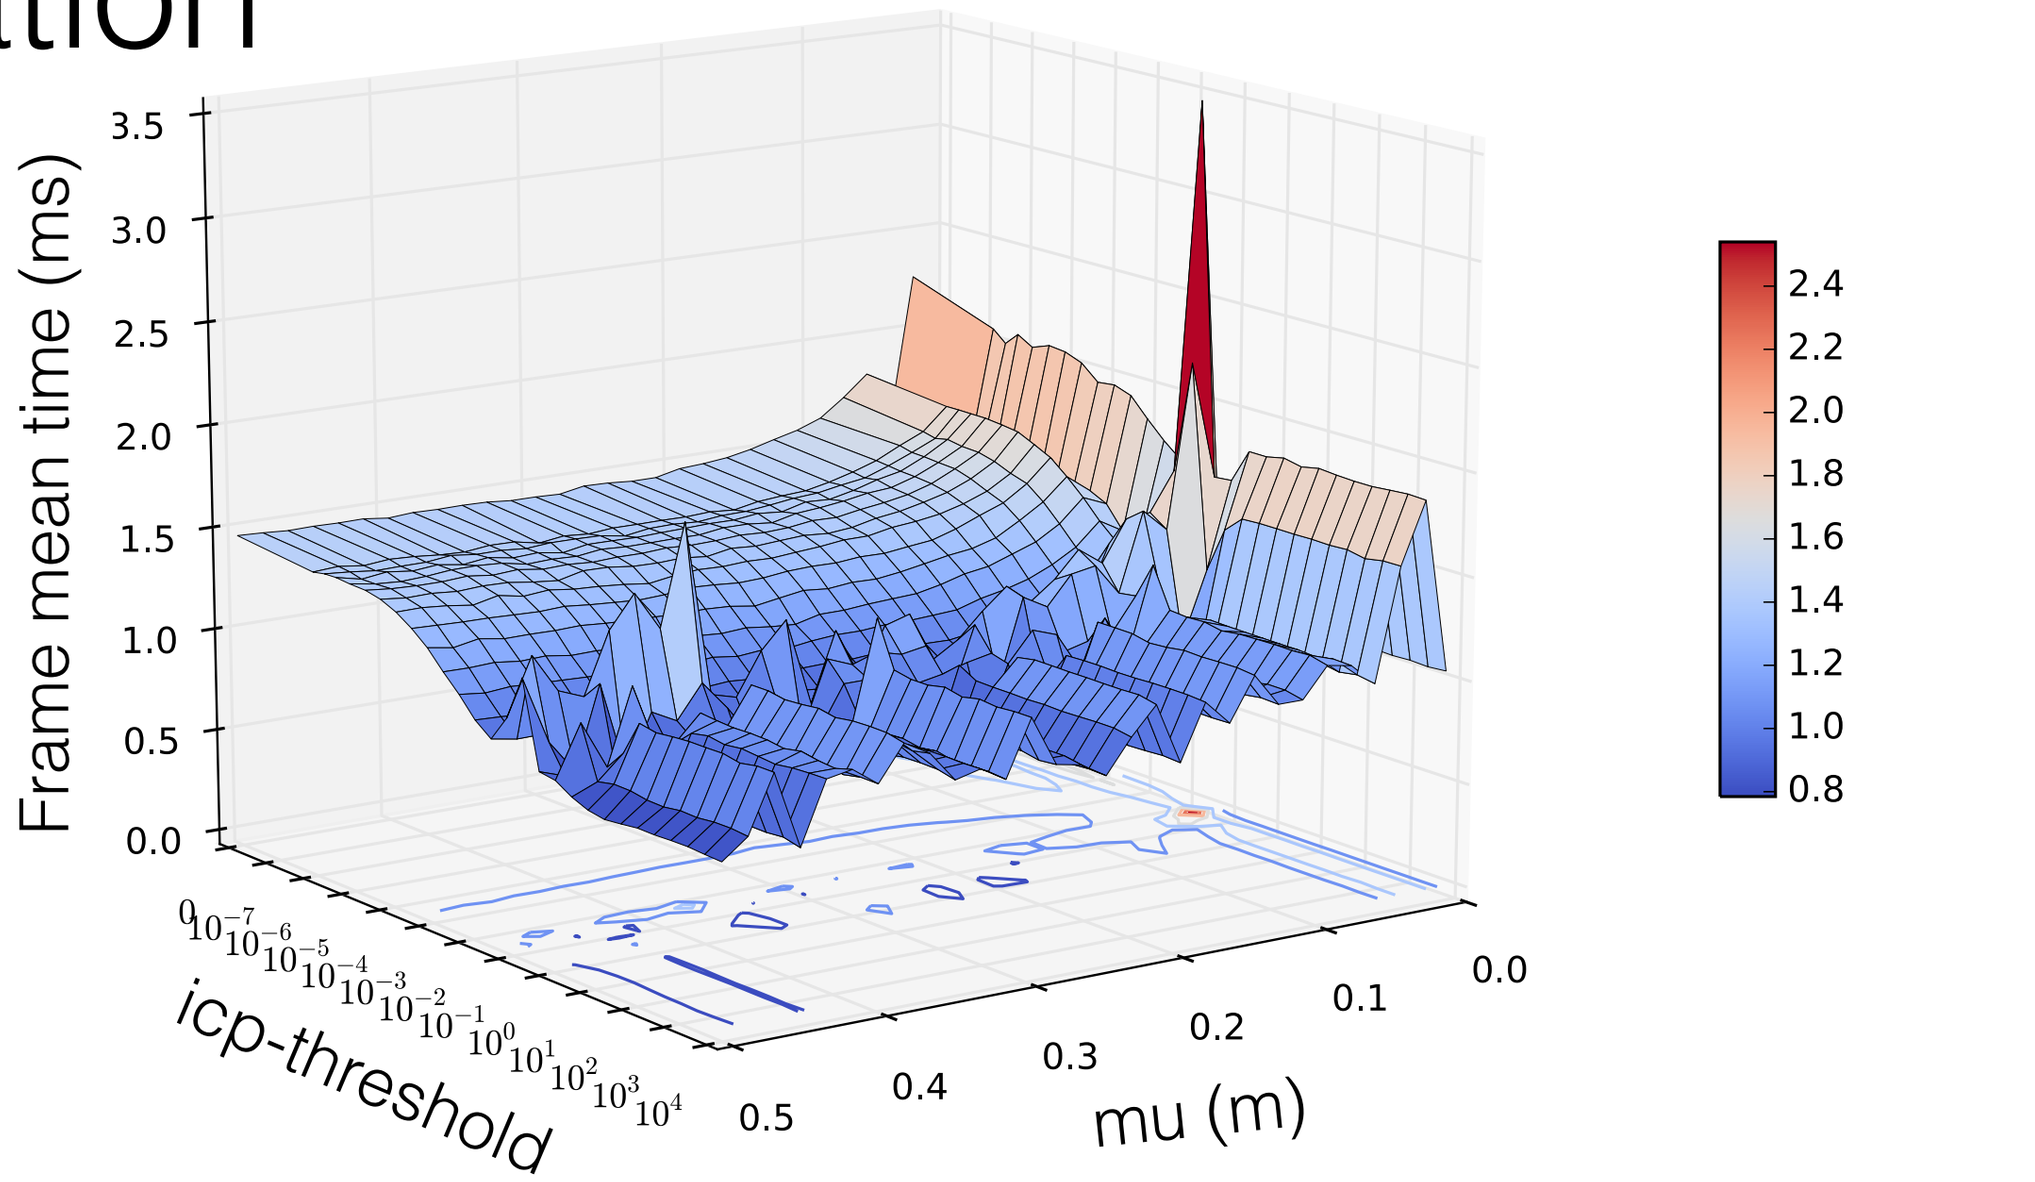

- KinectFusion runtime response surface: non-linear, multi-modal and non-smooth
- Optimal **algorithm configurability** enables better performance and better accuracy of the computation

Integrating Algorithmic Parameters into Benchmarking and Design Space Exploration in 3D Scene Understanding **(PACT 2016)**

Exploration goal illustration

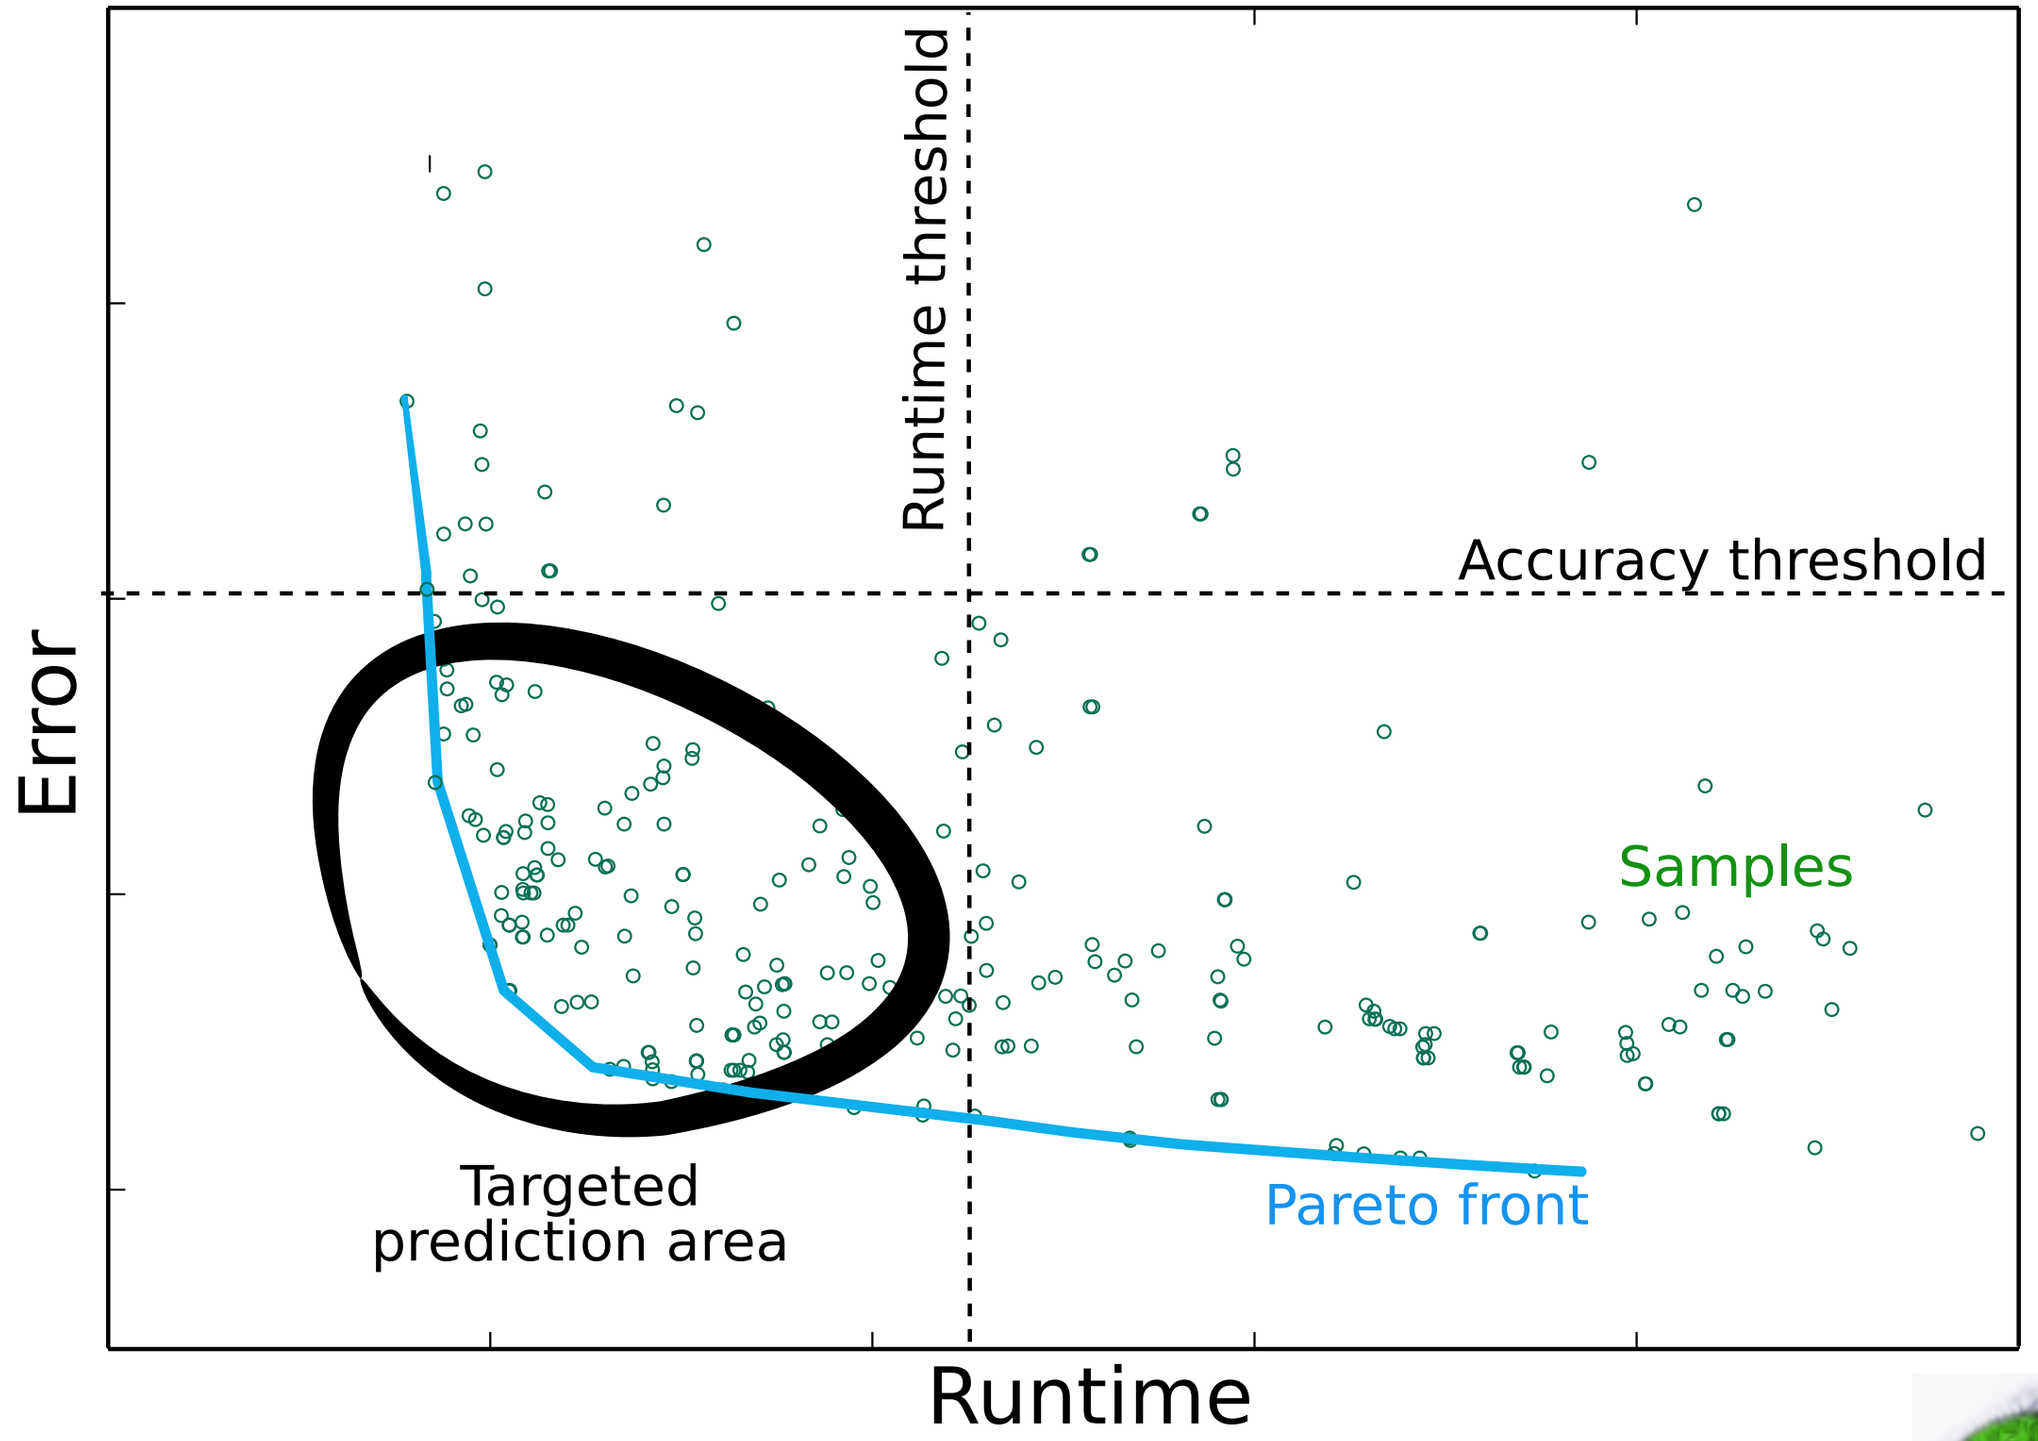

Algo Design-Space Exploration (DSE): Active Learning Methodology

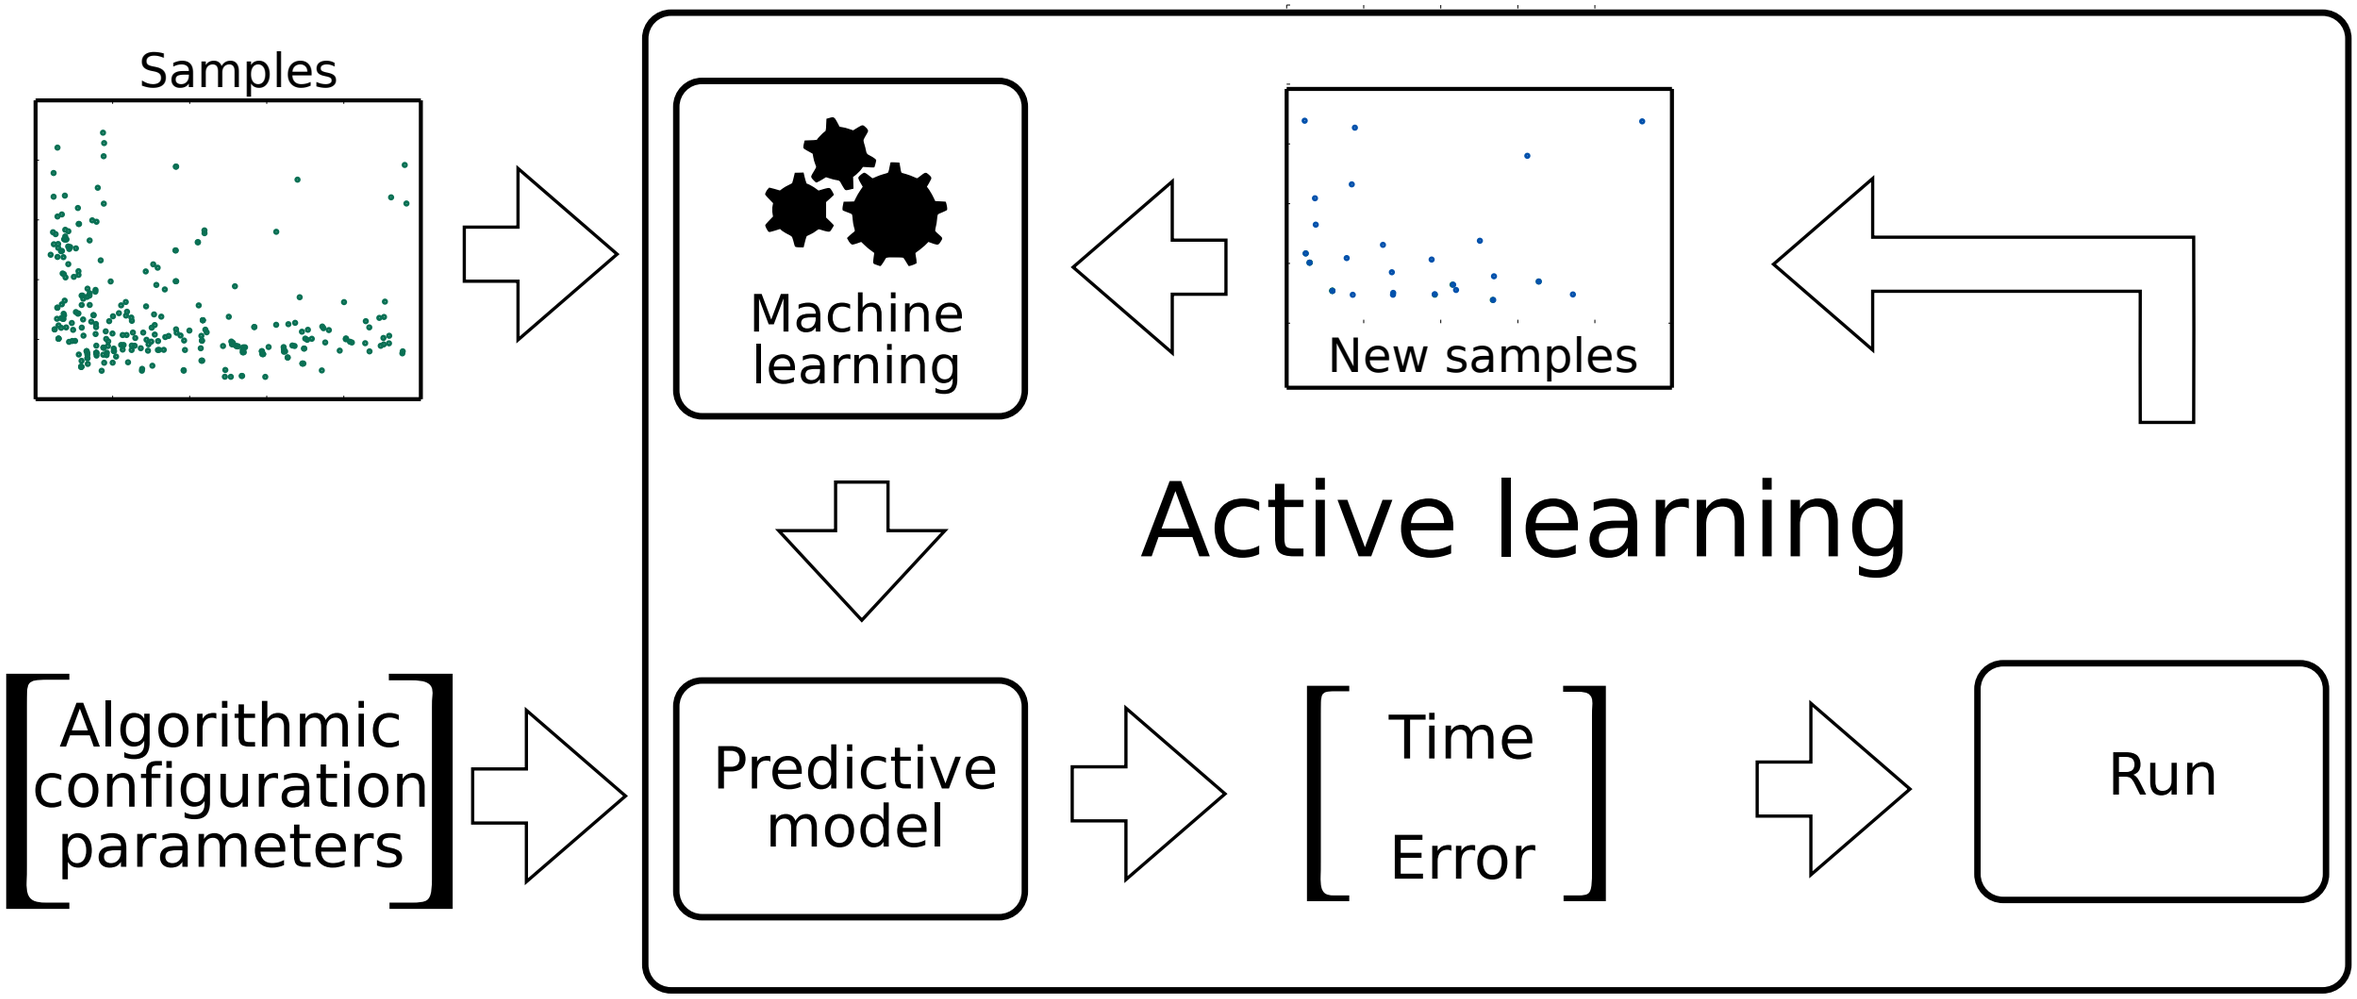

Machine learning methods used

Decision Tree

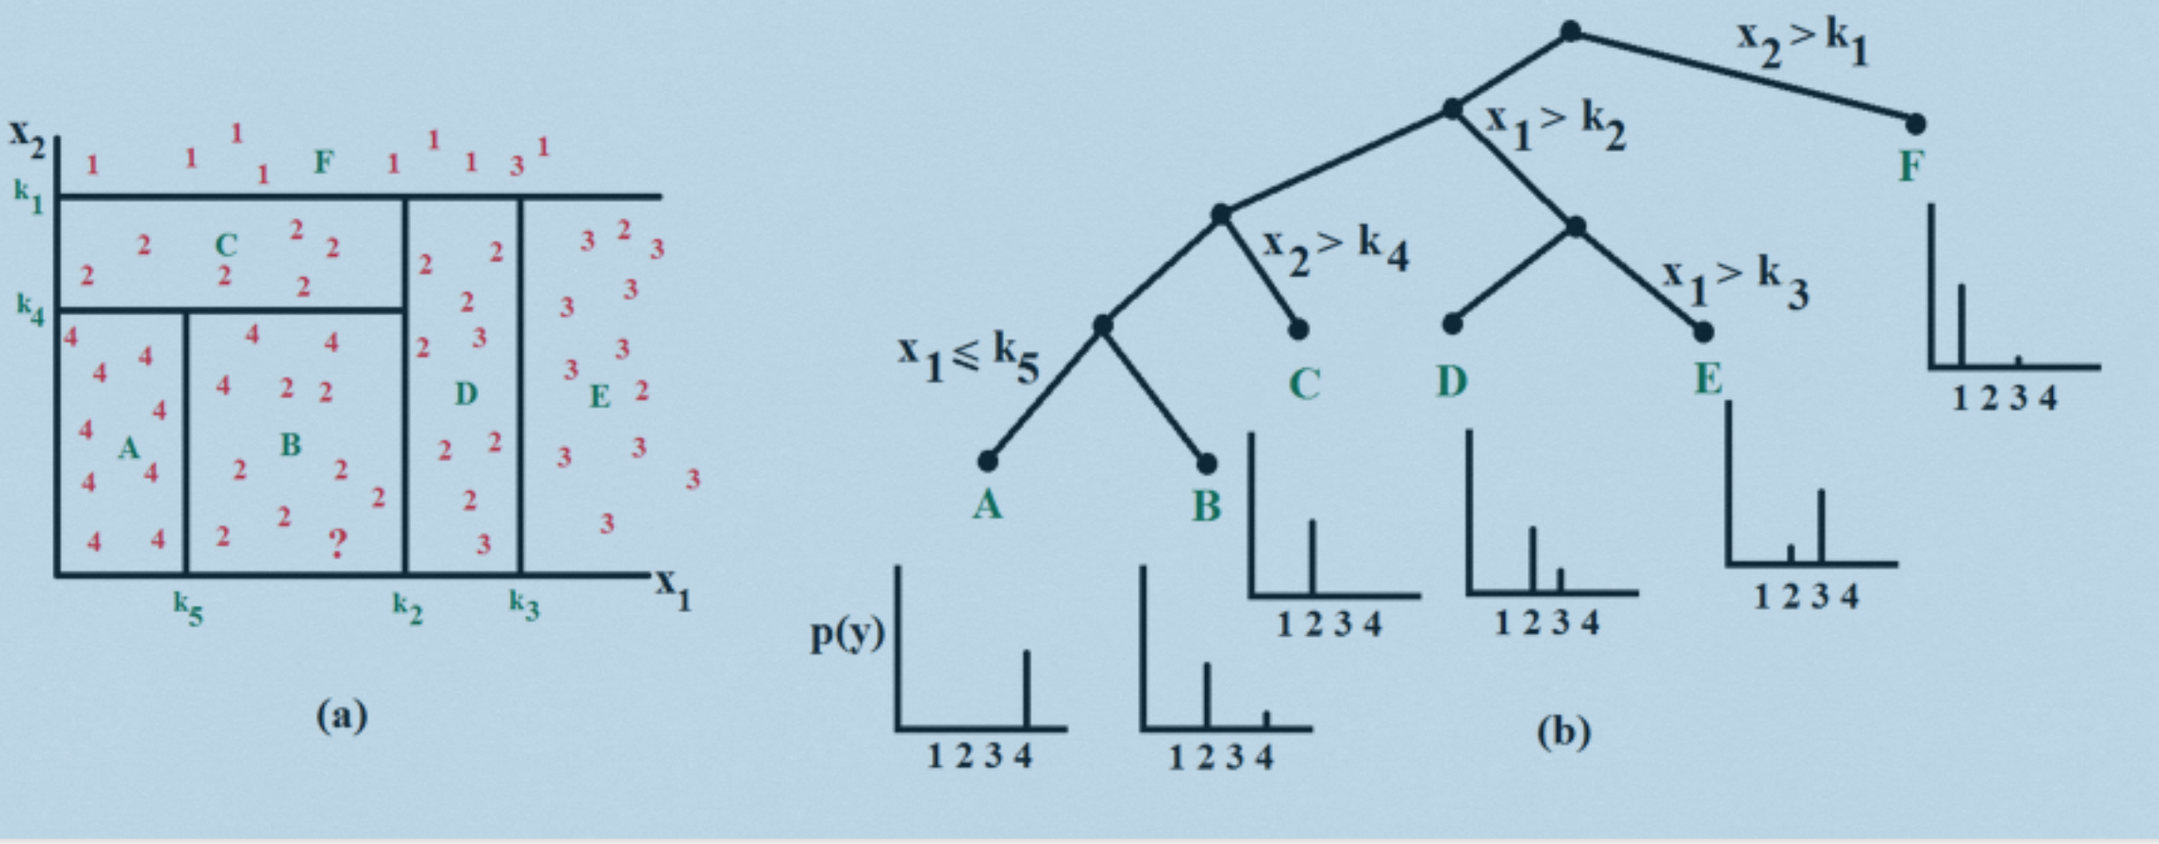

Random Forest

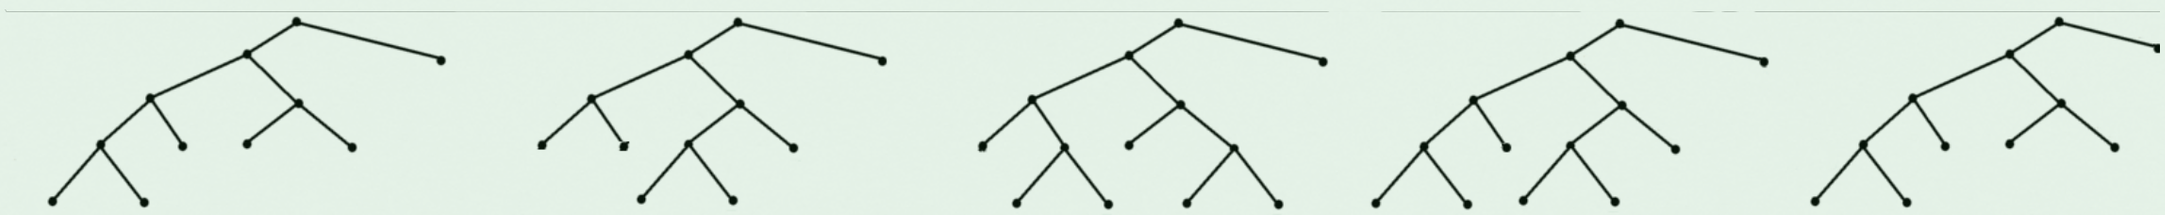

Results **ElasticFusion** DSE error/runtime III

Machine	Type	CPU	CPU name	CPU cores	GPU	GPU name
NVIDIA/Intel	Desktop	Intel Ivy Bridge	E5-1620	8	NVIDIA	GTX 780 Ti

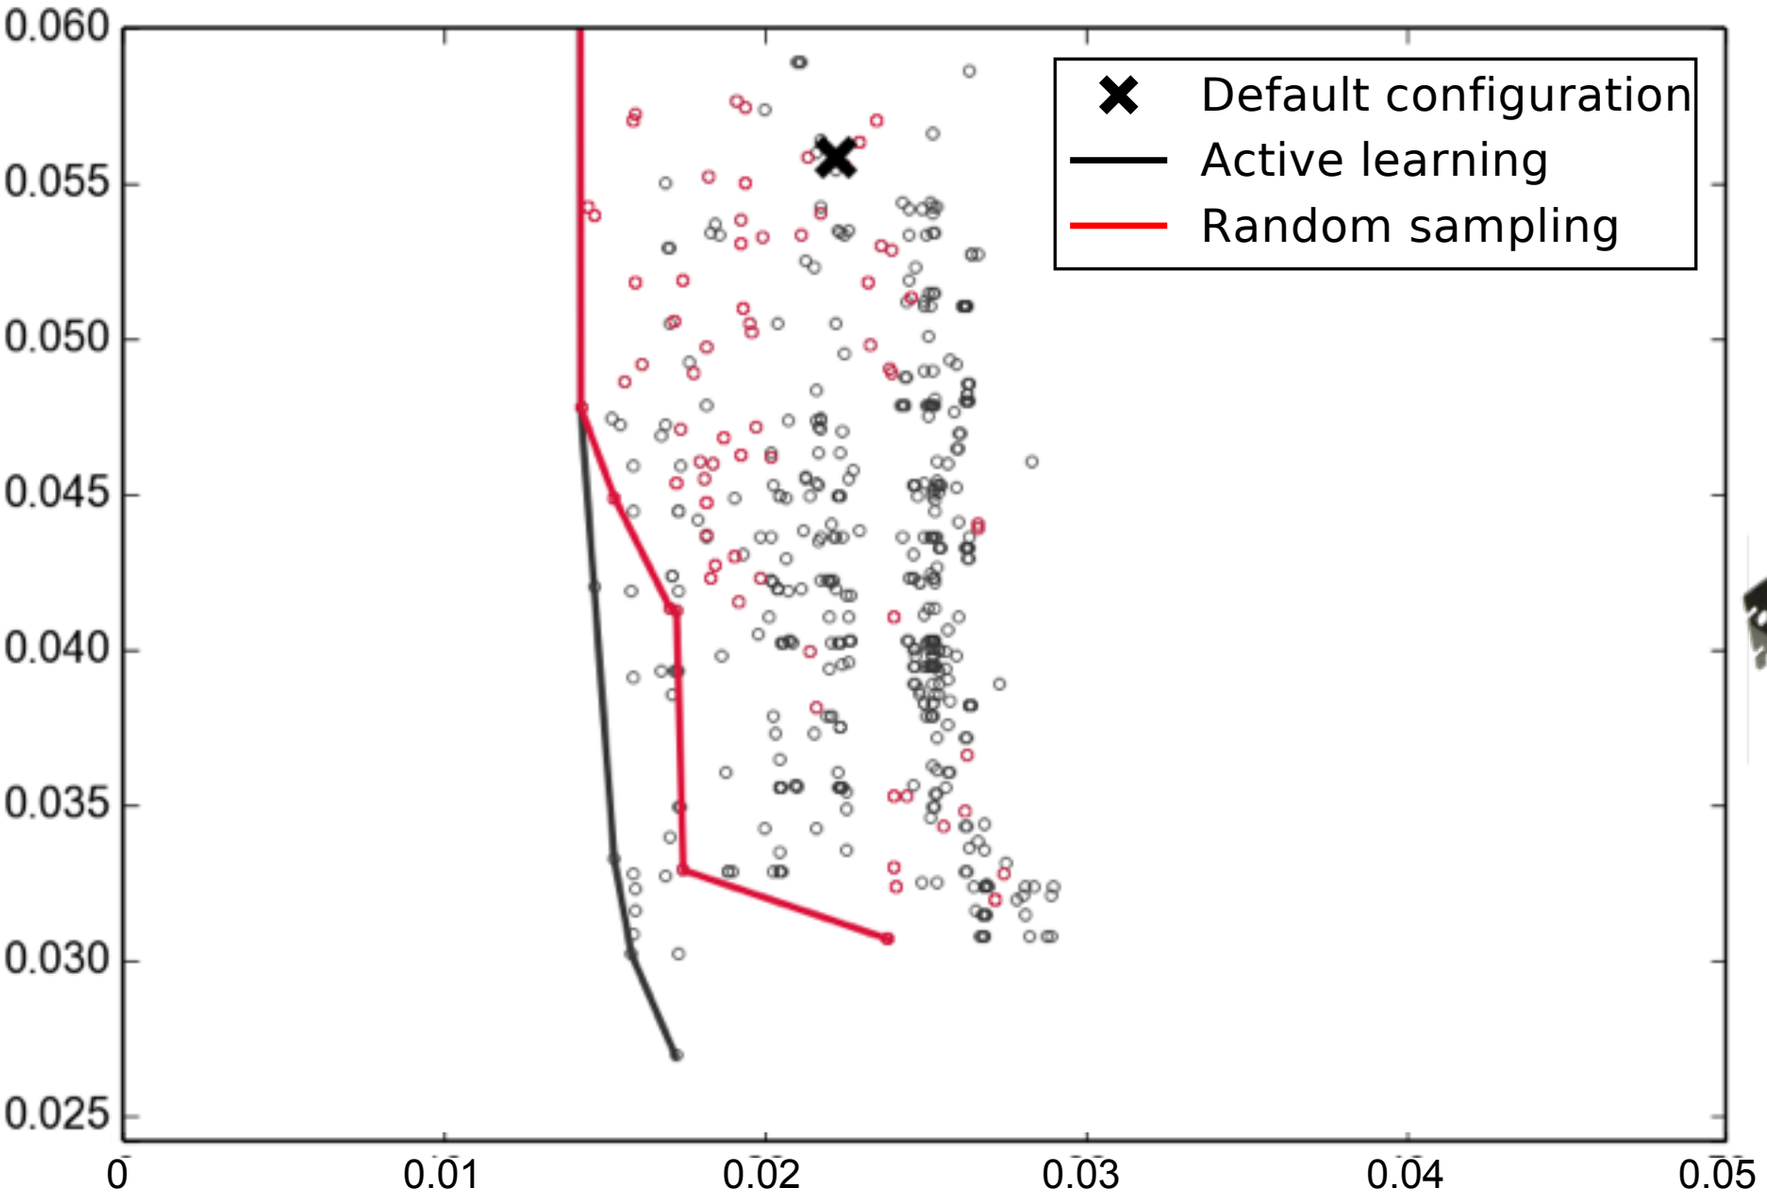

Results **ElasticFusion** DSE error/runtime III

Machine	Type	CPU	CPU name	CPU cores	GPU	GPU name
NVIDIA/Intel	Desktop	Intel Ivy Bridge	E5-1620	8	NVIDIA	GTX 780 Ti

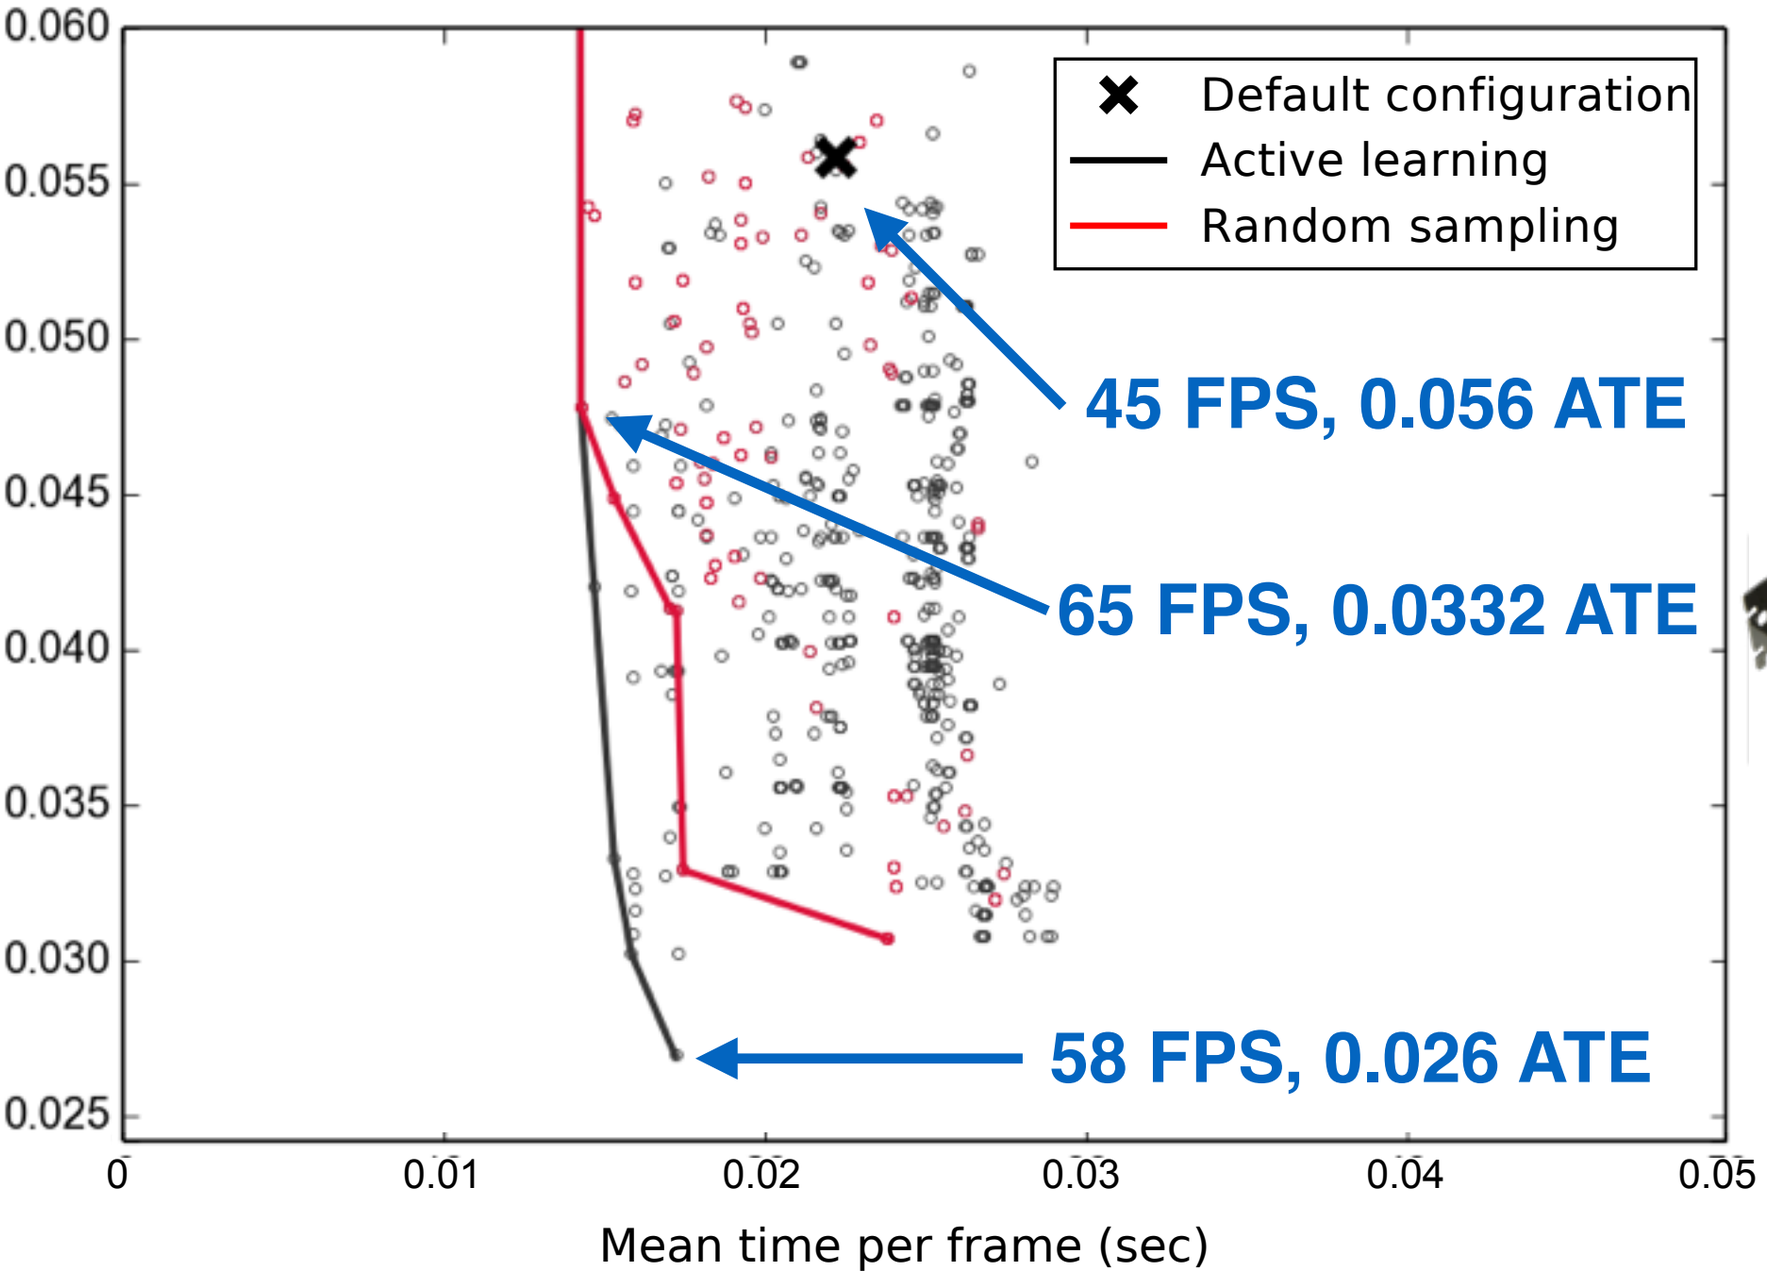

Luigi Nardi - Imperial College London

Results KinectFusion DSE error/runtime I

Machine	Type	CPU	CPU name	CPU cores	GPU	GPU name
Hardkernel **ODROID-XU3**	Embedded	ARM A15 + A7	Exynos 5422	4 + 4	ARM	Mali-T628

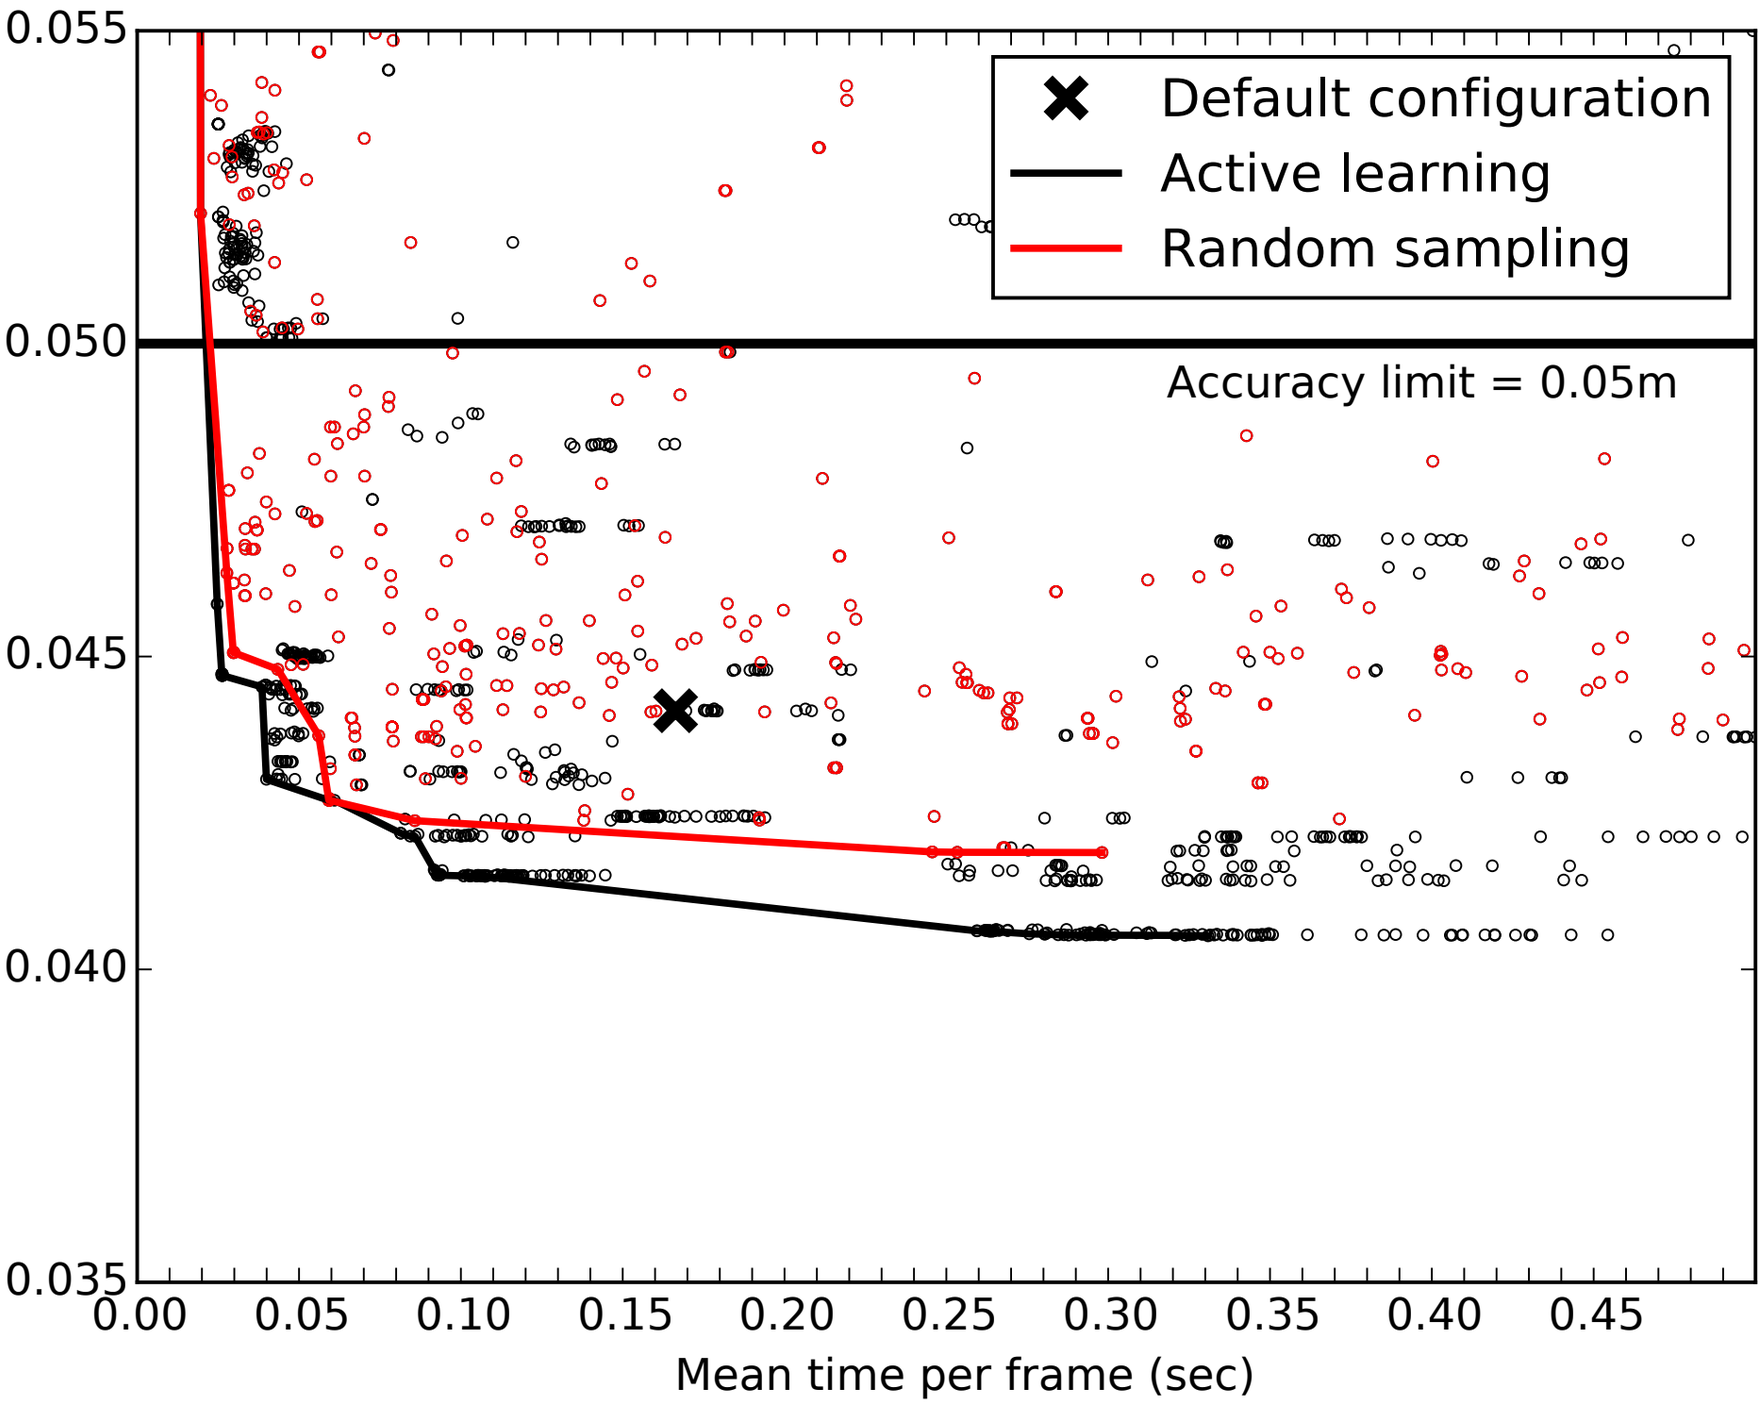

Luigi Nardi - Imperial College London

Results KinectFusion DSE error/runtime I

Machine	Type	CPU	CPU name	CPU cores	GPU	GPU name
Hardkernel **ODROID-XU3**	Embedded	ARM A15 + A7	Exynos 5422	4 + 4	ARM	Mali-T628

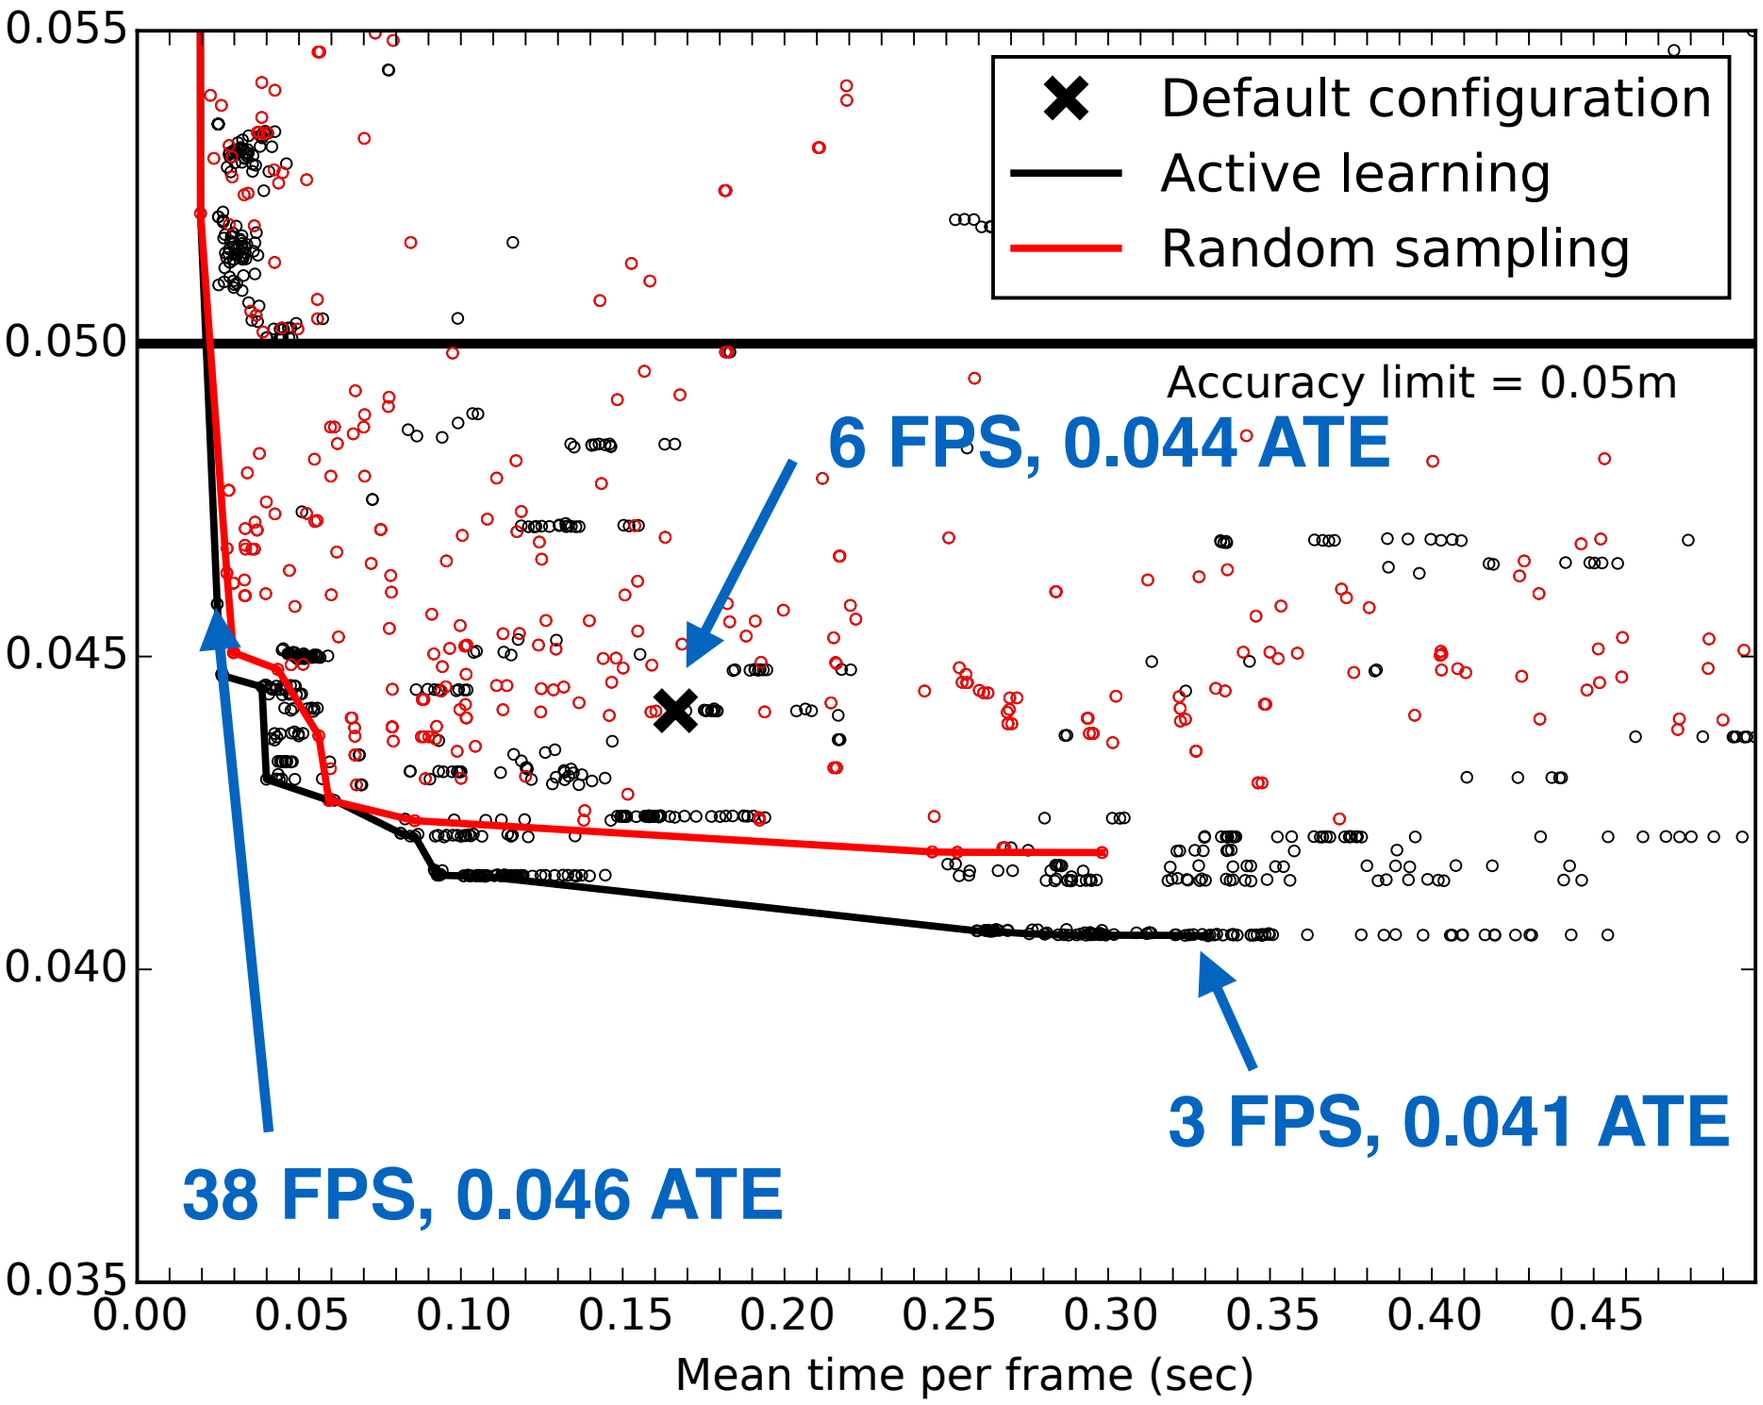

Results KinectFusion DSE error/runtime II

Machine	Type	CPU	CPU name	CPU cores	GPU	GPU name
ASUS T200TA	Detachable laptop	Intel Silvermont	Atom Z3795	4	Intel	HD Graphics

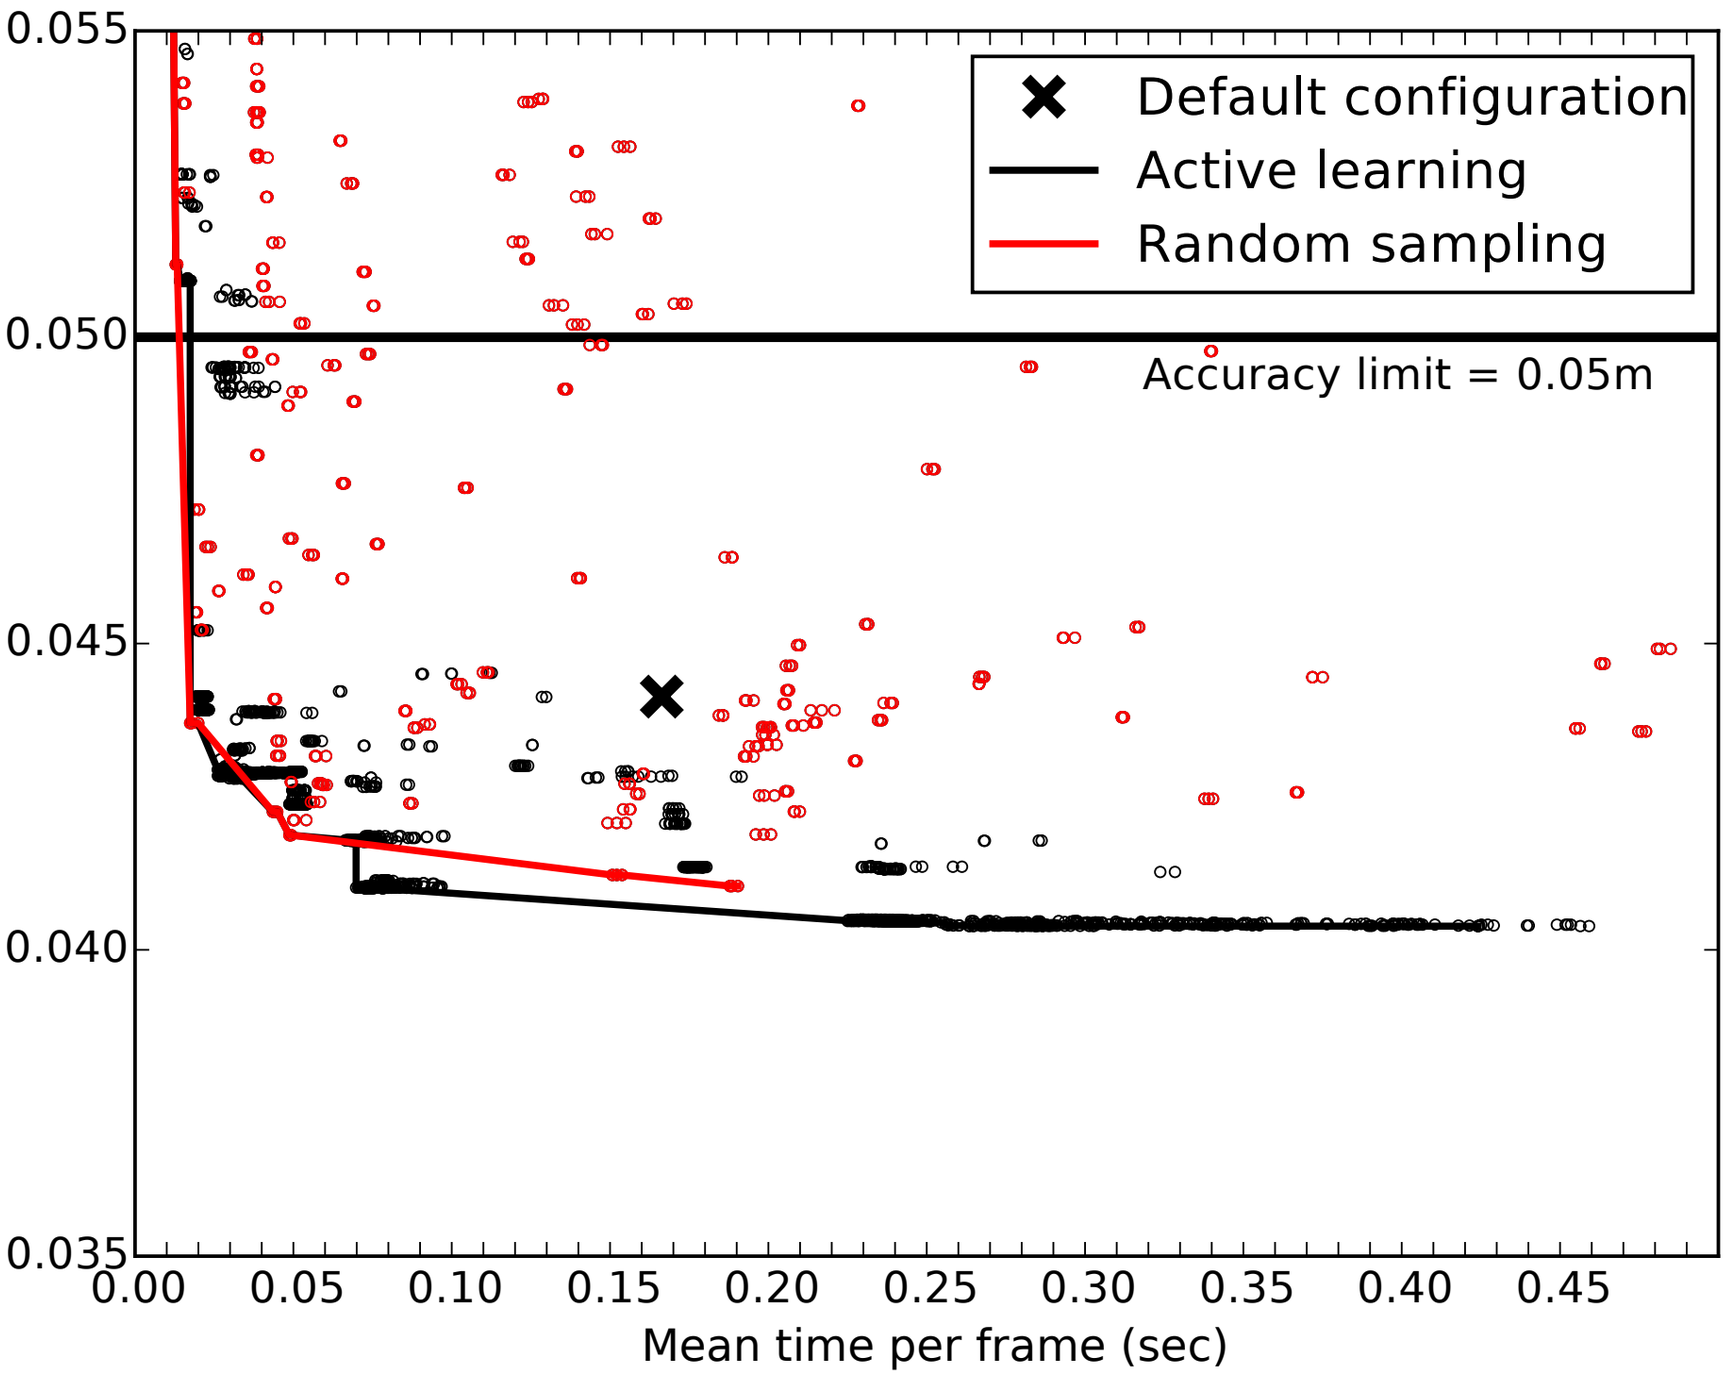

Conclusion - take away messages

1. Building tools to explore the performance landscape for SLAM solutions
2. Pareto maps how configurations should be adapted when objectives change - static and dynamic
3. Performance and accuracy improvement over default
4. Generalisation to other applications

Live demos GTC-EU 2016

Imperial College London

Booth D3

Passenger Terminal Hall 1/Main deck

September 28th and 29th

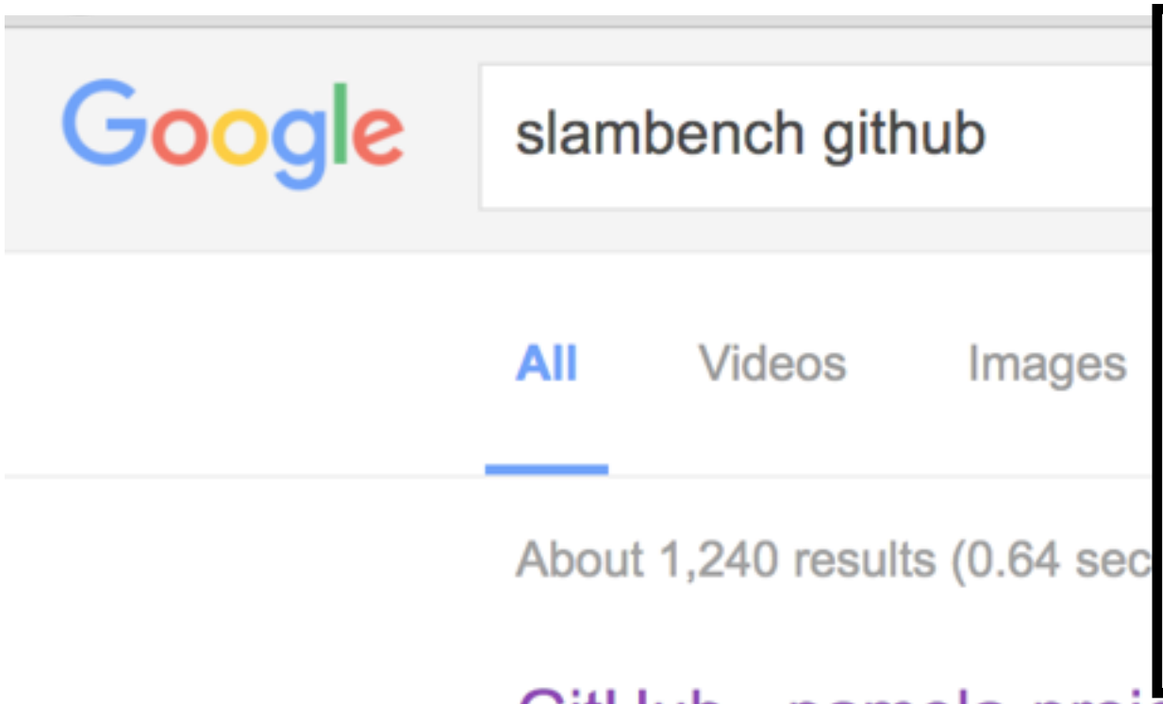

SLAMBench install
Booth D3 for help

Publicly released 13/11/2014
(1400+ downloads)

References I

- [Nardi et al. 2015] L. Nardi, B. Bodin, M. Z. Zia, J. Mawer, A. Nisbet, P. H. J. Kelly, A. J. Davison, M. Luján, M. F. P. O'Boyle, G. Riley, N. Topham, and S. Furber. "Introducing SLAMBench, a performance and accuracy benchmarking methodology for SLAM." Submitted, arXiv:1410.2167, 2015.
- [Newcombe et al. ICCV 2011] R. A. Newcombe, S. J. Lovegrove and A. J. Davison. "DTAM: Dense tracking and mapping in real-time." Computer Vision (ICCV), 2011 IEEE International Conference on. IEEE, 2011.
- [Rusinkiewicz and Levoy 2001] S. Rusinkiewicz, and M. Levoy. "Efficient variants of the ICP algorithm." 3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on. IEEE, 2001.
- [Chen et al. 2013] J. Chen, D. Bautembach, and S. Izadi, Scalable real-time volumetric surface reconstruction, in ACM Trans. Graph., 2013.
- [Newcombe et al. ISMAR 2011] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon. "KinectFusion: Real-time dense surface mapping and tracking." 10th IEEE Int. Symp. on Mixed and augmented reality (ISMAR), 2011.
- [Handa et al. 2014] A. Handa, T. Whelan, J. McDonald, and A. J. Davison. A Benchmark for RGB-D Visual Odometry, 3D Reconstruction and SLAM. IEEE Int. Conf. on Robotics and Automation, ICRA 2014.
- [Reitmayr] G. Reitmayr. KFusion github 2011. https://github.com/GerhardR/kfusion
- [Curless and Levoy 1996] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In Proc. Computer graphics and interactive technique. ACM, 1996.
- [Whelan et al. 2012] T. Whelan, M. Kaess, M. Fallon,H. Johannsson, J. Leonard, and J. McDonald. Kintinuous: Spatially extended kinectfusion. 2012.
- C. Jiawen, D. Bautembach, and S. Izadi. "Scalable real-time volumetric surface reconstruction." ACM TOG, 2013.
- Frahm, Jan-Michael, et al. "Building Rome on a cloudless day." Computer Vision–ECCV 2010. Springer Berlin Heidelberg, 2010.
- Erhan, Dumitru, et al. "Scalable object detection using deep neural networks." Proceedings of the IEEE CVPR. 2014.

References II

- Arbelaez, Pablo, et al. "Contour detection and hierarchical image segmentation." IEEE Pattern Analysis and Machine Intelligence, 2011.
- [Ogilvie 2014] Ogilvie, William, et al. "Fast automatic heuristic construction using active learning." Proceedings of the Workshop on Languages and Compilers for Parallel Computing (LCPC'14). 2014.
- [Siegmund 2015] Siegmund Norbert et al. "Performance-influence models for highly configurable systems", submitted FSE 2015.
- [Guo 2013] Guo, Jianmei, et al. "Variability-aware performance prediction: A statistical learning approach." Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on. IEEE, 2013.
- [Grewe 2011] Grewe, Dominik et al. "A static task partitioning approach for heterogeneous systems using OpenCL." Compiler Construction. Springer Berlin Heidelberg, 2011.
- [Kurek 2013] Kurek, Maciej, Tianchi Liu, and Wayne Luk. "MULTI-OBJECTIVE SELF-OPTIMIZATION OF RECONFIGURABLE DESIGNS WITH MACHINE LEARNING." 2nd Workshop on Self-Awareness in Reconfigurable Computing Systems (SRCS'13). 2013.
- [Balaprakash 2013] Balaprakash, Prasanna, Robert B. Gramacy, and Stefan M. Wild. "Active-learning-based surrogate models for empirical performance tuning." Cluster Computing (CLUSTER), 2013 IEEE International Conference on. IEEE, 2013.
- [Vespa 2015] Vespa Emanuele. "Sparse voxelization of dense volumetric reconstruction with automated analysis of scene reconstruction quality." M.Res. thesis, Imperial College London, 2015.

Copyrights

- Author: unknown. Microsoft Kinect camera. [Image]. Retrieved from http://channel9.msdn.com/Series/KinectSDKQuickstarts/Understanding-Kinect-Hardware
- Author: Dyson Ltd. Dyson 360 Eye. [Video]. Retrieved from https://www.youtube.com/watch?v=OadhuICDAjk
- Author: Google Inc. Google Tango project. [Image]. Retrieved from http://blogthinkbig.com/en/project-tango-googles-mobile-kinect/
- Author: unknown. Audi autonomous car. [Photograph]. Retrieved from http://www.wired.com/2010/06/audis-robotic-car-looks-hot-in-old-school-livery/
- Author: ExtremeTech. Google Shaft robot. [Photograph]. Retrieved from http://www.extremetech.com/extreme/173318-google-wins-darpas-robotics-challenge-wonders-if-it-was-a-good-idea-to-turn-down-future-military-contracts
- Author: HardKernel. ODROID-XU3 board. [Photograph]. Retrieved from http://www.hardkernel.com/main/products/prdt_info.php?g_code=G135235611947
- Author: PC Specialist Ltd. Vortex series laptop. [Photograph]. Retrieved from https://www.pcspecialist.co.uk/forums/showthread.php?23366-My-new-beast-15-6-quot-Vortex-III
- Author: Arndale.org. Arndale board. [Photograph]. Retrieved from http://www.arndaleboard.org/wiki/index.php/Main_Page
- Author: Unknown. Chip. [Image]. Retrieved from https://cajalesygalileos.wordpress.com/2013/06/23/un-chip-ultrasensible-identifica-15-cepas-de-gripe/
- Author: Unknown. Eye. [Image]. Retrieved from http://gallery.digitalculture.asu.edu/?/interactive-environments/computer-vision/
- Author: Unknown. Compiler. [Image]. Retrieved from http://d3q6qq2zt8nhwv.cloudfront.net/107/large-icon.png