CrossMark
← click for updates

STUDY PROTOCOL

Applying machine learning to automated segmentation of head and neck tumour volumes and organs at risk on radiotherapy planning CT and MRI scans [version 1; referees: 1 approved with reservations]

Carlton Chu[1], Jeffrey De Fauw[1], Nenad Tomasev[1], Bernardino Romera Paredes[1], Cían Hughes[1], Joseph Ledsam[1], Trevor Back[1], Hugh Montgomery[2], Geraint Rees[3], Rosalind Raine[4], Kevin Sullivan[5], Syed Moinuddin[5], Derek D'Souza[5], Olaf Ronneberger[1], Ruheena Mendes[5], Julien Cornebise[1]

[1]DeepMind, London, EC4A 3TW, UK
[2]Institute of Sport, Exercise and Health, London, W1T 7HA, UK
[3]Alexandra House, University College London, Bloomsbury Campus, London, WC1N 3AR, UK
[4]Department of Applied Health Research, University College London, London, WC1E 7HB, UK
[5]University College Hospital, London, NW1 2BU, UK

Abstract

Radiotherapy is one of the main ways head and neck cancers are treated; radiation is used to kill cancerous cells and prevent their recurrence.

Complex treatment planning is required to ensure that enough radiation is given to the tumour, and little to other sensitive structures (known as organs at risk) such as the eyes and nerves which might otherwise be damaged. This is especially difficult in the head and neck, where multiple at-risk structures often lie in extremely close proximity to the tumour. It can take radiotherapy experts four hours or more to pick out the important areas on planning scans (known as segmentation).

This research will focus on applying machine learning algorithms to automatic segmentation of head and neck planning computed tomography (CT) and magnetic resonance imaging (MRI) scans at University College London Hospital NHS Foundation Trust patients. Through analysis of the images used in radiotherapy DeepMind Health will investigate improvements in efficiency of cancer treatment pathways.

Corresponding author: Joseph Ledsam (jledsam@google.com)

Background

Cancers of the head and neck account for 2% of all cancers, and worldwide they account for around 300,000 deaths each year. Between 1990 and 2006 the incidence of different head and neck cancers has altered dramatically. The incidence of oral cavity cancer has risen by more than 30% to 3.02 per 100,000 and oropharyngeal cancer incidence has more than doubled with a significant change in causation (human papilloma virus rather than smoking or alcohol) and at-risk subpopulation (younger rather than older patients) (NCIN, 2012; OCIU, 2010; Parkin, 2010).

Most such cancers are treated using radiotherapy (CRUK, 2016). Planning such treatment involves delineating the tumour to be irradiated, and also the structures (organs at risk, OAR) to which administration of radiation should be minimised. The complex anatomy of the head and neck, where multiple OARs often lie in extremely close proximity to the tumour, make this process (known as 'segmentation') difficult: it can take radiotherapy experts four hours or more to do this (Harari et al., 2010). Furthermore as tumour and body shape change over a course of treatment (which can last weeks), it can be necessary to repeat the segmentation analysis at both fiscal and temporal cost; longer times between cancer diagnosis and treatment increase mortality and worsen outcomes (Chen et al., 2008; Mikeljevic et al., 2004).

Advances in machine learning have allowed the creation of sophisticated image recognition tools which might perform this process faster and at least as accurately. Machine learning has undergone a revolutionary transformation with the resounding success of so-called Deep Learning algorithms introduced by (Hinton et al., 2006) and demonstrated at scale by others (Krizhevsky et al., 2012). Those algorithms combine artificial neural networks, well-known since their introduction in the mid-50s (Rosenblatt, 1959), with advances in computational power and algorithmics which have enabled remarkable success in handling the deluge of high-dimensional "big data" (Krizhevsky et al., 2012). This resulted in the development and rapid deployment of automatic feature extraction, i.e. combining parts of the data into meaningful elementary units for higher-level processing, that would have previously required the painstaking trial-and-error process of manual design by a human.

Such processes might be readily applied to automated segmentation for radiotherapy planning in the treatment of head and neck tumours. The prevalence of head and neck cancers, and the complexity of radiotherapy planning, make them ideal targets for such an automated computer-based approach. An automated segmentation system would allow planning to start immediately after a patient is scanned. Such a service could increase speed to treatment. It could also help reduce variation in radiotherapy outcomes between centres (Peters et al., 2010) by standardising planning prior to dose simulation while increasing the efficiency of patient workflow. Although many techniques have been proposed for automatic segmentation in radiotherapy (Daisne & Blumhofer, 2013; Rohlfing et al., 2005), none have shown sufficiently good performance for routine use in clinical care.

In order for machine learning algorithms to reach expert levels at image segmentation, they must first learn from existing data. This study aims to achieve this with a dataset of expertly labelled images from previously-treated patients at University College London Hospital (UCLH) NHS Foundation Trust (London, UK), with the ultimate objective of improving outcomes for patients with head and neck cancers.

Aims and objectives

Primary objective

1.1
To investigate the feasibility of developing computer algorithms that can identify important anatomical structures and the cancers themselves in head and neck cancer planning scans to help target radiotherapy treatment.

Secondary objectives

Should the primary objective be accomplished, we intend to validate performance using retrospective data through:

2.1
Assessment of quality of automated segmentation using retrospective planning CT images. Expert radiation oncologists, blind to image source, will assess both automatically and manually segmented scans, and determine whether segmentation met a standard for clinical use.

Study design

This is a retrospective, non-interventional study. Analyses performed in the study will be on fully anonymised medical images (computed tomography (CT) and magnetic resonance imaging (MRI) scans, labelled with manual segmentation, dose threshold and cancer type).

The protocol follows similar procedures to De Fauw et al. (2016).

Inclusion criteria

Patients who received radiotherapy treatment for head and neck cancers at UCLH NHS Foundation Trust between 01/01/2008 and 20/03/2016 will be eligible for inclusion in this study.

Exclusion criteria

Data from patients who have previously manually requested that that their data should not be shared, even for research purposes in anonymised form, and have informed the UCLH NHS Foundation Trust of this, will be ineligible and removed by UCLH NHS Foundation Trust staff before research begins.

Sample size

Approximately 700 retrospective patient cases will be part of this study.

Most recent machine learning algorithms benefit from large datasets on which to train (tens to hundreds of thousands of data instances (Silver *et al.*, 2016)). Across all machine learning applications the predictive power (as percentage of data instances correctly classified) of the algorithm depends on the size and quality of the dataset.

The sample size is informed by the existing literature (Mnih, 2015; Silver *et al.*, 2016) and by DeepMinds previous work in the field of machine learning. We believe that the research goals are possible despite the relatively small number of scans, as compared to other research projects because of low variation between the different biological images and by limiting the scope of the research to segmentation rather than diagnosis.

Data
For all patients meeting inclusion/exclusion criteria the following electronic health record data will be required to complete this project successfully:

(1) CT scan(s) taken during the course of radiotherapy planning and treatment

(2) MRI scan(s) taken during the course of radiotherapy planning and treatment

(3) CT labelling information outlining anatomical and tumour volumes, with associated radiotherapy dose thresholds

(4) Information on what type of tumour is present in each image

(5) Patient gender and age group (to the nearest 5 years)

The anonymisation procedures adopted will remove any information not specified to further avoid transfer of patient identifiable information. All anonymisation will be formally verified by UCLH NHS Foundation Trust staff before transfer.

Algorithm development
In order to develop the algorithms, DeepMind will work with the labelled medical image files to apply machine learning and AI techniques including but not limited to: supervised and semi-supervised convolutional neural networks, recurrent neural networks, unsupervised clustering, reinforcement learning (Murphy, 2012).

Statistical analysis
Descriptive statistics (such as the Dice similarity coefficient, average surface distance and maximal surface distance) will be used to compare the quality of algorithm segmentation against the expert reference segmentation during algorithm training.

In order to assess the accuracy of the final model segmentation a retrospective test subset of radiotherapy planning images will also have both manual and automatic segmentations corrected by expert clinicians who have not seen the images before and who are blinded to how the segmentation was produced. The same statistical methods described above will be used to compare the ground truth manual segmentation, automatic segmentation and clinically corrected manual and automatic segmentations.

Data protection
Anonymisation
This study requires existing retrospective data only; no prospective data are needed nor will be collected from patients, hospitals or healthcare workers. No direct patient contact will occur and necessary data will be anonymised from this source dataset.

Anonymisation of all image files and clinical information is performed and validated at UCLH NHS Foundation Trust before transfer. No patient identifiers will be transferred to DeepMind. In addition the data will be protected to HSCIC Information Governance standards and access is strictly controlled to prevent any attempt to re-identify the data.

Data storage
DeepMind Health has developed and established a state-of-the-art secure patient information handling service utilising Common Criteria EAL4 compliant firewalls and on-disk encryption (using Advanced Encryption Standard with a 256-bit key) of all research data, all housed within an ISO 27001 compliant data centre. After anonymisation data will be transferred to our London, UK data centre. This data handling facility conforms to NHS HSCIC Information Governance Statement of Compliance Toolkit (formally assessed at level 3).

Access will be granted by the custodian of data and no other members of the team. Only those working directly on the data in a research capacity will have access.

Data destruction
The data sharing agreement between DeepMind and UCLH NHS Foundation Trust lasts for 5 years. After this period the agreement will be reviewed should future work seek to build on this project. After the data sharing agreement expires all data used in the study will be destroyed. No modification will be made based on the data after destruction.

Data destruction will involve the deletion of the encryption/decryption keys for all project volumes, and 8-pass random data write to all physical disks within the DeepMind Health data infrastructure. A certificate of destruction will be provided to the Trust.

The algorithms developed during the study will not be destroyed. DeepMind Health knows of no way to recreate the patient images transferred from the algorithms developed. No patient identifiable data will be included in the algorithms.

Ethical considerations
Ethical approvals
The research on the dataset received formal Research Ethics Committee approval on 6th April 2016 (REC reference 16/SC/0189).

Consent
No patient will be approached directly. Only anonymised retrospective data collected as part of routine clinical care are included. In such cases the ICO code of practice states that explicit consent is not generally required (ICO, 2012).

Adverse events

The project is non-interventional and does not involve any direct patient contact. All patient data is historical and all patients have completed their radiotherapy treatment prior to data transfer.

Monitoring

The study will be monitored both internally and externally. Internally DeepMind managers (TB, JL) will oversee and monitor progress on a day-to-day basis, ensuring the protocol is adhered to and no compliance issues arise.

Clinical and methodological experts (RM, DD, KS, SAM, GR, RR) are working with DeepMind to further oversee the ethical, clinical and methodological considerations of the project and will advise on at least a weekly basis to ensure no deviation from the described protocol.

Externally the information governance team at the UCLH NHS Foundation Trust will be consulted before commencing data collection.

Access

DeepMind has access to the required data to support the research aims of this study. To ensure compliance with the common law principle of data confidentiality, DeepMind will only receive anonymised data from UCLH NHS Foundation Trust. DeepMind works with the Trust to ensure accuracy and clarity in the data to allow useful and consistent interpretation at all times.

Dissemination

The results will be disseminated through normal academic channels, initially focusing on conference proceedings and the indexed peer reviewed literature relevant to the fields of machine learning, artificial intelligence, radiotherapy and clinical research.

DeepMind will engage in patient and public involvement groups during the research study.

Conclusion

We propose an exploratory study covers an initial testing of machine learning algorithms for automatic segmentation of head and neck planning CT and MRI scans. The results will be assessed against expert segmentation.

Author contributions

All authors contributed to study design and methodology. CC, CH, JD, NT, BP, OR and JC contributed to machine learning approaches. RM, KS, DD and SAM contributed expertise in oncology, radiation physics and radiography. TB and JL contributed to project steering and information governance. GR, RR and HM contributed to methodological oversight.

Competing interests

University College London Hospital NHS Foundation Trust administration time spent on this work will be paid to the Trust by DeepMind.

The Chief Investigator (CC) and some co-investigators (TB, JC, CH, OR, JD, NT, BP) are full time employees of DeepMind. GR, RR, HM and JL are paid consultants of DeepMind. The company is also funding the research.

References

Cancer Research UK: **Oral Cancer Statistics.** Cancer Research UK. 2016; (last accessed 16/05/2016).
Reference Source

Chen Z, King W, Pearcey R, *et al.*: **The relationship between waiting time for radiotherapy and clinical outcomes: A systematic review of the literature.** *Radiother Oncol.* 2008; **87**(1): 3–16.
PubMed Abstract | **Publisher Full Text**

Daisne JF, Blumhofer A: **Atlas-Based automatic segmentation of head and neck organs at risk and nodal target volumes: a clinical validation.** *Radiat Oncol.* 2013; **8**: 154.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

De Fauw J, Keane P, Tomasev N, *et al.*: **Automated analysis of retinal imaging using machine learning techniques for computer vision [version 1; referees: 1 approved].** *F1000Research.* 2016; **5**: 1573.
Publisher Full Text

Harari PM, Song S, Tomé WA: **Emphasizing conformal avoidance versus target definition for IMRT planning in head-and-neck cancer.** *Int J Radiat Oncol Biol Phys.* 2010; **77**(3): 950–958.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Hinton GE, Osindero S, Teh YW: **A Fast Learning Algorithm for Deep Belief Nets.** *Neural Comput.* 2006; **18**(7): 1527–54.
PubMed Abstract | **Publisher Full Text**

Information Commissioner's Office: **Anonymisation: managing data protection risk code of practice.** 2012; (last accessed 03/06/2016).
Reference Source

Krizhevsky A, Sutskever I, Hinton GE: **ImageNet Classification with Deep Convolutional Neural Networks.** In *NIPS.* 2012; **1**: 4.
Reference Source

Mikeljevic JS, Haward R, Johnston C, *et al.*: **Trends in postoperative radiotherapy delay and the effect on survival in breast cancer patients treated with conservation surgery.** *Br J Cancer.* 2004; **90**(7): 1343–1348.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Mnih VM, Hinton GE: **Learning to Detect Roads in High-Resolution Aerial Images.** *European Conference on Computer Vision.* 2010; **6316**: 210–223.
Publisher Full Text

NCIN: **Potentially HPV-related head and neck cancers. NCIN Data Briefing.** 2012.
Reference Source

Oxford Cancer Intelligence Unit: **Profile of Head and Neck Cancers in England:**

Incidence, Mortality and Survival. 2010.
Reference Source

Parkin DM, Boyd L, Walker LC: **16. The fraction of cancer attributable to lifestyle and environmental factors in the UK in 2010.** *Br J Cancer.* 2011; **105**(Suppl 2): S77–81.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Peters LJ, O'Sullivan B, Giralt J, *et al.*: **Critical impact of radiotherapy protocol compliance and quality in the treatment of advanced head and neck cancer: results from TROG 02.02.** *J Clin Oncol.* 2010; **28**(18): 2996–3001.
PubMed Abstract | **Publisher Full Text**

Rohlfing T, Brandt R, Menzel R, *et al.*: *Quo vadis*, **atlas-based segmentation?**

In *The Handbook of Biomedical Image Analysis: Segmentation and Registration Models.* Edited by Suri J, Wilson DL, Laxminarayan S. New York, NY: Kluwer Academic/Plenum Publishers. 2005; 435–486.
Publisher Full Text

Rosenblatt F: **The perceptron: a probabilistic model for information storage and organization in the brain.** *Psychol Rev.* 1958; **65**(6): 386–408.
PubMed Abstract | **Publisher Full Text**

Silver D, Huang A, Maddison CJ, *et al.*: **Mastering the game of Go with deep neural networks and tree search.** *Nature.* 2016; **529**(7587): 484–489.
PubMed Abstract | **Publisher Full Text**

Open Peer Review

Current Referee Status: ?

Version 1

Referee Report 31 October 2016

? **Yuan Feng**

Soochow University, Suzhou, China

This study protocol presented the procedure to acquire and analyse medical image data for automatic segmentation. The protocol is clear besides the following points need to be clarified:

1. The authors presented the algorithm used for segmentation very briefly. This part should be elaborated to inform the details of each algorithm used. Also, citation of Murphy 2012 was not found in the reference.

2. The statistical analysis was not well presented. The evaluation and statistical methods should be described in detail.

3. There are many different segmentation studies in the literature, some specifically aimed for radiotherapy should be cited.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Competing Interests: No competing interests were disclosed.