

WestminsterResearch

<http://www.westminster.ac.uk/westminsterresearch>

Competing for Pixels: A Self-Play Algorithm for Weakly-Supervised Semantic Segmentation

Shaheer U. Saeed, Shiqi Huang, João Ramalhinho, Iani J.M.B. Gayo, Nina Montaña-Brown, Ester Bonmati, Stephen P. Pereira, Brian Davidson, Dean C. Barratt, Matthew J. Clarkson and Yipeng Hu

This is a copy of the author's accepted version of a paper subsequently published in IEEE Transactions on Pattern Analysis and Machine Intelligence, 47 (2), pp. 825-839.

It is available online at:

<https://doi.org/10.1109/TPAMI.2024.3474094>

© 2025 IEEE . This manuscript version is made available under the CC-BY 4.0 license

<https://creativecommons.org/licenses/by/4.0/>

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners.

Competing for pixels: a self-play algorithm for weakly-supervised semantic segmentation

Shaheer U. Saeed, Shiqi Huang, João Ramalinho, Iani J.M.B. Gayo, Nina Montaña-Brown, Ester Bonmati, Stephen P. Pereira, Brian Davidson, Dean C. Barratt, Matthew J. Clarkson, Yipeng Hu

Abstract—Weakly-supervised semantic segmentation (WSSS) methods, reliant on image-level labels indicating object presence, lack explicit correspondence between labels and regions of interest (ROIs), posing a significant challenge. Despite this, WSSS methods have attracted attention due to their much lower annotation costs compared to fully-supervised segmentation. Leveraging reinforcement learning (RL) self-play, we propose a novel WSSS method that gamifies image segmentation of a ROI. We formulate segmentation as a competition between two agents that compete to select ROI-containing patches until exhaustion of all such patches. The score at each time-step, used to compute the reward for agent training, represents likelihood of object presence within the selection, determined by an object presence detector pre-trained using only image-level binary classification labels of object presence. Additionally, we propose a game termination condition that can be called by either side upon exhaustion of all ROI-containing patches, followed by the selection of a final patch from each. Upon termination, the agent is incentivised if ROI-containing patches are exhausted or disincentivised if a ROI-containing patch is found by the competitor. This competitive setup ensures minimisation of over- or under-segmentation, a common problem with WSSS methods. Extensive experimentation across four datasets demonstrates significant performance improvements over recent state-of-the-art methods. Code: <https://github.com/s-sd/spurl/tree/main/wss>

Index Terms—Self-Play, Weak Supervision, Segmentation.

I. INTRODUCTION

Segmentation is a fundamental task in computer vision with demonstrated applications in e.g., medical image processing [67], [76], autonomous driving [19], and robotics [63]. Fully supervised segmentation systems, trained using human-annotated data, have widely been proposed, however, they are limited by the expensive pixel-level annotation acquisition for datasets of increasing sizes [63]. Tasks or datasets, where annotation cost is high or where these labels are not readily available, are especially hindered by this requirement [63] e.g., in the medical domain where expert time may be restricted. Weakly supervised semantic segmentation (WSSS) aims to

S.U. Saeed, S. Huang, J. Ramalinho, I.J.M.B. Gayo, N. Montaña-Brown, E. Bonmati, D.C. Barratt, M.J. Clarkson and Y. Hu are with the UCL Hawkes Institute, Department of Medical Physics and Biomedical Engineering, University College London, London WC1E 6BT, U.K.

S. Huang is also with the School of Optics and Photonics, Beijing Institute of Technology, Beijing, China.

E. Bonmati is also with the School of Computer Science and Engineering, University of Westminster, London W1W 6UW, U.K.

S.P. Pereira is with the Institute for Liver and Digestive Health, University College London, London NW3 2QG, U.K.

B. Davidson is with the Division of Surgery and Interventional Sciences, University College London, London WC1E 6BT, U.K.

E-mail: shaheer.saeed.17@ucl.ac.uk

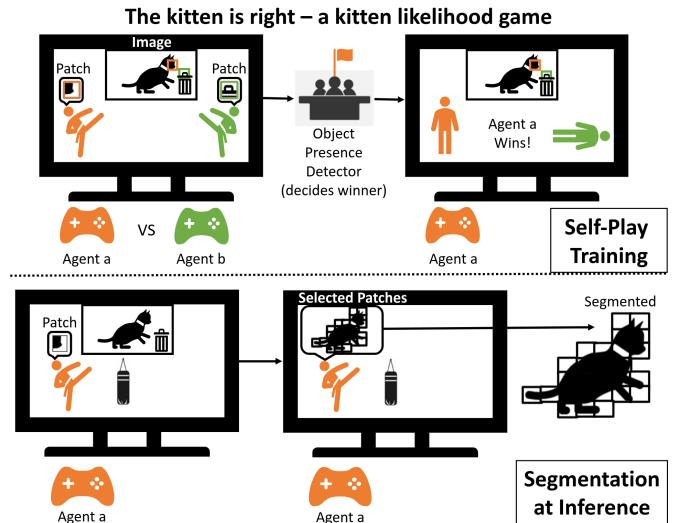


Fig. 1. Competitive self-play for segmentation using only weak signals generated by an object presence detector (itself trained using only image-level classification labels). At inference the trained agent plays against a dummy opponent to segment the image.

directly address these challenges by utilising weak labels for training, such as image-level classification labels (i.e., binary labels of object presence) [112].

We propose a WSSS method in which we gamify segmentation such that two opposing agents compete to segment a region-of-interest (ROI), where each move is scored by a weak signal derived from a pre-trained object presence detector. This detector is a classifier of object presence, trained only using image-level classification labels. This weak signal is similar to classifiers used in recent WSSS methods [112] and ensures discriminative features, able to distinguish object existence from non-existence, are learnt only using image-level labels. The framework is based on reinforcement learning (RL) self-play where the competitive setup encourages minimisation of over- or under-segmentation, a common problem with WSSS [11], [111].

In our framework, an image is first divided into patches using a grid of a pre-defined size, where the patches are akin to playing-cards in a card game. An agent from each competing side (two sides in this case) selects, without replacement, a patch that it deems most likely to contain a part of a ROI. The game is turn-based and thus prevents selection of the same patch by opposing sides. Each patch is scored by the pre-trained object presence detector, using the unbounded

classification probability, or logits, and the agent with a higher score wins the turn and is given a positive reward signal. This is repeated until one of the agents delivers a termination signal, which indicates when an agent considers all ROI-containing patches to be exhausted. The final patches from each side are then scored by the object presence detector and if the patch from the side that delivered the terminal signal is scored higher, then that side wins the game - leading to a high positive reward signal. The main idea is summarised in Fig. 1. This may be similar to an innate self-play-style approach within the human brain, where object presence likelihood for parts of the field-of-view (FOV) may be decided before localisation [103].

Learning to maximise the score/ reward using RL, in this competitive setup, ensures that under-segmentation is minimised as it is not beneficial to terminate the game early. This is because early termination may mean that the opponent could potentially select a more object-like patch, or the opponent's selected patch may be scored higher simply based on chance e.g., if one part of the object is more object-like than the other. It also ensures that over-segmentation is minimised, as timely termination is rewarded, once ROI-containing patches are exhausted. If the game is not terminated appropriately and both sides keep selecting patches with no ROIs, the opposing side may win the game based on chance as one part of the background may be more object-like compared to another. Over thousands of games, these early or later termination cases, lead to reduced expectation for the reward, compared to an effective termination strategy which terminates only when ROI-containing patches are exhausted. Thus RL, which learns by maximising the expectation of reward learns an effective termination strategy. In practice, limiting numbers of turns (by discouraging termination beyond the specified range) and selecting patch sizes sufficiently smaller than approximate ROI size were sufficient constraints to learn effective segmentation. Once trained, an agent may select patches without opposition to get a segmentation with resolution equal to patch-size.

Using self-play not only improves performance, by ensuring that over- or under-segmentation are minimised through an effective termination strategy, but also proves more time-efficient compared to scoring all patches using the object presence detector as non-selected patches do not need to pass through the detector. It also avoids the need for empirical thresholds for classification probability for patch-based segmentation. If pixel-level segmentation is required, the grid can be shifted over by half the patch-size and the game repeated, where the final segmentation may be obtained using a majority vote. This still proves more time-efficient compared to a sliding windows approach where the object-presence detector must score all patches.

A. Summary of contributions

1) We propose a gamified competitive self-play framework for image segmentation which minimises over- or under-segmentation, a common problem for other WSSS methods; 2) We propose to train this WSSS framework using RL where the reward is based solely on an object presence detector trained only using image-level classification labels of object presence;

3) We evaluate our proposed approach using two benchmarks datasets for WSSS (PASCAL VOC 2012 and MS COCO 2014) as well as two real-world medical datasets, representing real-world applications where pixel-level human annotations are expensive to acquire and curate (liver tumour and prostate gland segmentation).

B. Summary of motivation

Over- and under-segmentation of ROIs are known common problems with WSSS methods [11], [111] (refer to literature review in Sec. II-B for further details). CAM- and MIL-based techniques, reliant on object classifiers, are prone to under segment due to classifiers activating only highly-discriminative object features as opposed to entire objects, where these features may be under-specific for full object segmentation. Other common WSSS approaches, based on erasure of highly discriminative features so classifiers can activate other less discriminative features [12], [30], [41], [42], [45], [84], [93], [94], [106], aim to overcome these problems but are prone to over-segment with a high sensitivity to hyperparameters, often even requiring per-image tuning, such as the the number of times of erasure, area of erasure etc. Methods based on an aggregation of tasks or images for learning discriminative features face similar problems with high sensitivity to hyperparameters to control the aggregation (e.g., number of tasks or images to use, how to aggregate features etc.).

We propose to gamify segmentation in a competitive set-up where two agents compete to segment a ROI. The competitive set-up ensures that over- and under-segmentation are minimised by means of a termination condition that ensures all ROI-containing pixels/ patches of a given image are exhausted. We design a reward signal for training the competing agents, based on object presence likelihood in patches of an image where appropriate termination is rewarded. The object presence likelihood itself is learnt solely using image-level labels. The game termination can be triggered by either of the competing sides and, upon termination, each side presents a final patch, for which object presence likelihood is quantified. If the game is terminated too early by the agent (under-segmentation), the competitor can select an object-containing patch with a higher object presence likelihood, and thus a negative reward is delivered to the agent that called the termination. This negative reward minimises under-segmentation. If the game is terminated too late (over-segmentation), and all object-containing patches are already exhausted, then both competitors are likely to select patches with no object within. This may mean that, by chance, the opposing side can select a patch that has a higher object presence likelihood than the agent that called termination. The higher object presence likelihood, for the competitor's patch, in this case is likely to be based on random chance, with object presence likelihood for both patches being near-zero as all object-containing patches have already been exhausted. A negative reward is delivered to the agent that called the termination in this case as well. As some of these cases of over- or under-segmentation are based on random chance, they may not occur on every iteration. However, the average expectation for the reward for these over-

or under-segmented cases would be lower than an optimum termination, that leads to an optimum segmentation. We train the agents using RL, which maximises the expectation of reward and thus learns an appropriate termination strategy that minimises over- and under-segmentation.

Our approach differs from previous work by posing segmentation as an auto-competing process where the signals derived from the competition, and the competitive pressure by means of the designed termination-based reward, prove sufficient to minimise segmentation errors i.e., over- and under-segmentation, common with other WSSS methods.

II. RELATED WORK

A. Reinforcement learning self-play

Self-play describes a special class of multi-agent problems in RL where two or more agents compete to accomplish a task, often set-up as a zero-sum game [15], [78]. In self-play, an agent competes with itself, or a copy of itself, to improve without the need for any direct supervision e.g., human labels. Rather, a reward signal, often based on the win/loss conditions of a game, guides learning where both copies of the agent compete to maximise the reward, or in other words to win the game. This promotes learning competitive behaviours such as developing counters to effective opposing high-reward strategies or developing defensive strategies [15], [79], [80]. Perhaps, the main benefit of self-play comes from acquiring multiple sets of experiences that can be used to train a single agent [28], [29], [79], [80]. This not only promotes faster learning compared to learning against a hand-crafted strategy, but also allows higher diversity in experiences for potentially improved generalisability [3], [4], [28], [29], [79], [80]. Self-play also aids learning by allowing a near-perfect adversary, since the agent plays against an opponent that is equally skilled, a copy of itself, and improves as the opponent improves [44].

Thus far, the focus for self-play research has mostly been to use learning in the context of existing zero-sum games e.g., Chess [80], Go [79], Poker [29] and Hide and Seek [5]. These games and self-play research in general can be classified by information-completeness based on the comparative information accessible to each of the players [15], [79], [80]. Chess, where both players have the same information, and are fully aware of the opposition's payoffs and actions, are known as information-complete. Any games where an aspect of the game dynamics is hidden from any of the players are known as information-incomplete e.g., Hold'em Poker where the player is unaware of the observed state (card selection) for opposing players. Self-play has been extensively validated for both types of games.

Roots of self-play are grounded in game theory, however, RL self-play was popularised by works that proposed general algorithms that enabled learning in information-complete environments such as Go [79], [80] and Othello [86]. They showed, for the first time, that RL self-play starting from random play, could outperform handcrafted strategies and sophisticated search techniques without any domain knowledge, except game rules. These techniques were built upon and applied to

various complex information-complete computer games such as Tennis and Soccer [59].

Concurrently, research into using RL self-play for information-incomplete games also gained popularity through its use in games e.g., Poker [28], [29]. Further developments led to the application and extension of RL self-play to more challenging environments such as Capture the Flag [85], DouDizhu [104] and Robotic Soccer [23].

The goal with self-play applied to any game or competitive task is to reach a Nash equilibrium, whereby an agent will learn a perfect policy from which it will not deviate [4], [28], [29], [66], [75], [79], [80]. However, in practice, reaching Nash equilibrium is often not feasible due to the large action and observation spaces in modern games, unless human knowledge is used to abstract domains to manageable sizes [29]. Nonetheless self-play using RL (without injecting any expert knowledge) has managed to learn near-optimal solutions to classical information-complete games [79], [80]. Several opponent sampling curricula have been proposed for efficient learning in an attempt to get closer to Nash equilibria in more complex environments.

In 'vanilla self-play', often described as the 'best response' curriculum, the current most up-to-date version of the agent is used as an opponent. This opponent has been trained on the most experience against itself, and thus may offer the best response to an average response of the opposing player [29]. This insight prompted the use of a mixture of average response and best response as an opponent instead of only the best response, as using only marginally weaker opponents may promote learning [7], [29], [44], [49], [66], [88]. The key insight from practical experiments being that a new mixed curriculum may lead to more generalisable agents due to the larger breadth of collected experience and also to avoidance of cyclic learning and local optima common in vanilla self-play, where improvements are not made [7], [28], [49], [75], [88]. This method of learning with an opponent with mixed responses (mixture of best and average responses) is known as 'fictitious self-play'. In recent practice, in RL self-play, this may be implemented by randomly sampling past versions of agents (including the current agent itself) to act as opponents [59], [66], [88], [91]. This holds theoretically as an expectation of randomly sampling all past responses is equivalent to the average response. Any curriculum that uses non-uniform sampling distributions over all past versions of the agent may be classed as 'prioritised fictitious self-play' [88]. This often leads to faster learning, dependant on the sampling distribution, as only limited amounts of experience are collected against very weak opponents i.e., experience collection time is not wasted on much weaker opponents that pose no challenge to the current agent [88]. Developments in competitor sampling have enabled various real-world applications.

In recent years, self-play research has focused on finding competitive gamified formulations for a variety of problems (that are not pre-existing games) to find effective novel solutions in various domains. Kajiura et al. [37] demonstrated the use of RL self-play in image re-targeting by proposing a competition between agents to select optimal image operators, to adjust images to arbitrary sizes without information loss,

within reasonable times. Self-play found effective strategies, with processing times three orders of magnitude less than previous multi-operator methods. Hammar et al. [24] showed that by gamifying infrastructure security attacks e.g., denial of service and cross-site scripting, and defence, e.g., firewalls, effective intrusion prevention strategies, that outperform tested baselines, could be learnt with RL fictitious self-play without expert knowledge. Recently, Wang et al. [91] proposed RL prioritised fictitious self-play for protein engineering to mutate sequences towards desired properties, similar to playing pieces in a game, which led to the discovery of proteins with 7.8-fold bio-luminescence improvement compared to naturally occurring sequences.

B. Weakly supervised semantic segmentation

Instead of full pixel-level annotations, WSSS utilises weak labels to guide segmentation, e.g., image-level [2], [35], [94], patch-level [14], [47], [55], [64], [81] or pointers and sparse free-hand annotations [56], [57], [69], [87]. In this work, we focus on weak signals derived from image-level labels as they are often most time-efficient to obtain and are already available in many real-world applications e.g., medical imaging diagnostic tasks [112]. Furthermore, many works assume image-level labels as a pre-requisite for freehand-, or patch-based WSSS, because labelling e.g., patches or points, automatically leads to a binary annotation of object presence [112]. Missing image-level classification to ROI correspondence in WSSS, in ground truth data, remains a challenging task to overcome with previously proposed solutions being based mainly on class activation maps (CAM) [2], [35], [53], [94], or multi-instance learning (MIL) [34], [50], [96], [99].

Methods for WSSS based on CAM, extract ROIs from internal activations of a classification neural network, trained to classify object presence [74]. These methods, however, may be prone to under-segment since classifiers are driven to activate only small portions of the image with strong discriminative capabilities [11], [111]. Solutions employ a number of techniques. Region growing extrapolates the initial map, with or without external information [32], [40], [90]. The so-called erasure-based methods remove highly discriminative features from images such that classifiers activate other less discriminative features [12], [30], [41], [42], [45], [84], [93], [94], [106]; erasure area, number of erasures, are all hyperparameters to be tuned, often per image. Background modelling creates maps of the background with the help of external saliency maps [38], [48], [97], [100] or using a class-wise discriminator re-trained for every class [17]. Information aggregation accumulates information from multiple images to build object representations [18], [60], [83], [111]. Self-supervised learning uses auxiliary tasks to build supervisory signals [8], [11], [77], [90]. However, a general solution to WSSS using only image-level classification labels remains an active research area.

MIL has been proposed for WSSS due its ability to capture correspondence between bags of multiple instances, with available annotations, and single instances, to-be-labelled [21]. For WSSS, works aim to pose images as bags with image-level classification labels available and pixels as instances, to

be labelled [34], [50]. Some solutions define bags as patches with labels available [25]. Others incorporate clustering [99], or redefine instances as lattice patches to obtain coarse segmentation [96], which may increase inference time due to the need to infer for all patches. Due to a focus on discriminative features as in CAM, possibly sufficient for classification but under-specific for full-object segmentation, these formulations also tend to under-segment.

It should be noted that, more broadly, weakly supervised segmentation, dealing with guiding segmentation using image-level labels, may be categorised into semantic, instance and panoptic segmentation. While this work only focuses on semantic segmentation that assigns a class-label to each pixel in an image, other weakly supervised segmentation approaches use similar algorithmic approaches. Weakly supervised instance segmentation (WSIS) aims to distinguish between different instances of the same class within an image. Similar to WSSS, WSIS can use CAMs, where peaks in the activation can be used to distinguish between different object instances [33]. Often WSSS approaches can also be used to generate pseudo-labels to train WSIS systems [22], [39], [113]. Weakly supervised panoptic segmentation aims to solve both WSSS and WSIS simultaneously and assigns class labels to each pixel while simultaneously identifying different instances of objects, using largely similar algorithmic approaches to WSIS and WSSS [51].

This work focuses on WSSS and evaluates the approach for commonly used WSSS benchmarks.

III. METHODS

In this work, we propose a gamified formulation for WSSS where two agents compete to select ROI-containing patches until exhaustion of all such patches. The game score at each time-step is determined by an object presence detector which indicates the likelihood of a patch containing the ROI. The game score is used as a reward signal to train the agents using RL, and the object presence detector is pre-trained using only image-level object presence classification labels. The game may be terminated by either of the competing agents, once they consider all ROI-containing patches to be exhausted. If, post-termination, the opposing side is able to select a patch with higher likelihood of object presence (determined by the object presence detector), a negative reward signal is delivered to the side that called the termination signal, otherwise a positive reward is delivered. This competitive set-up ensures minimisation of over- or under-segmentation, as described in Sec. I and outlined in the following subsections.

This formulation involves two types of functions: 1) object presence detector - a classifier trained using binary image-level classification labels of object presence, to classify object presence within an image or patch; and 2) the agent - responsible for patch selection such that a segmentation map may be generated, trained using RL self-play with rewards derived from the trained object presence detector. We first formalise our proposed prioritised fictitious self-play and then describe the set-up for WSSS.

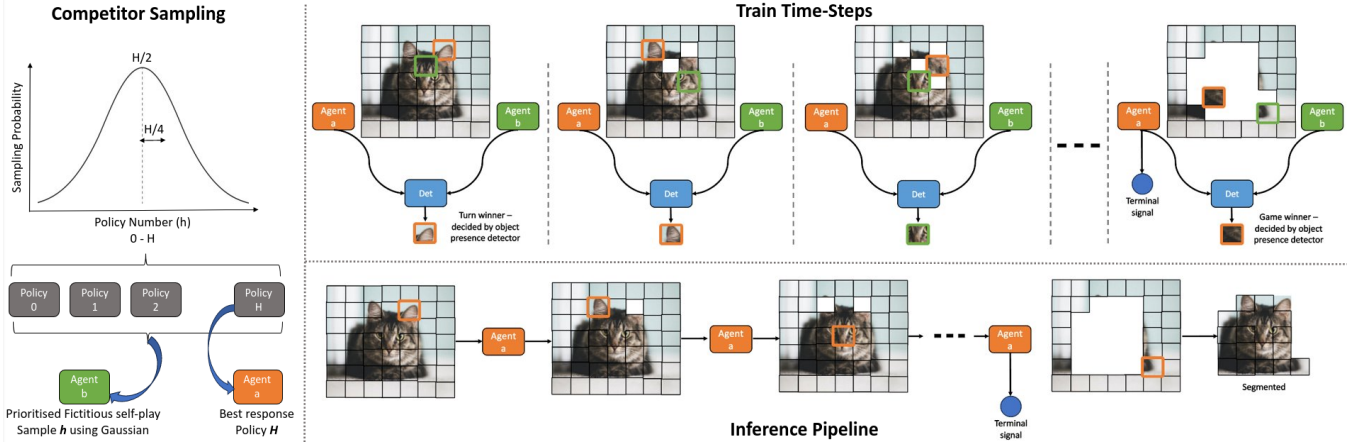


Fig. 2. Overview of the gamified WSSS segmentation using RL self-play.

A. Competing functions and the environment

In the form of self-play of interest in this work, two opposing agents compete to achieve a goal in a zero-sum game. The problem formulation is described in this section.

Single agent reinforcement learning: RL involves learning optimal policies for decision-making under Markov decision process (MDP) environments. The environment may be defined as $(\mathcal{S}, \mathcal{A}, p, r, \pi, \gamma)$.

States, actions and the policy: An agent or policy, interacting with an environment, from time-step t , $\pi(\cdot; \theta) : \mathcal{S} \times \mathcal{A} \rightarrow \{0, 1\}$ represents the probability of performing action $a_t \in \mathcal{A}$ to influence the environment, given observed state $s_t \in \mathcal{S}$, where \mathcal{A} and \mathcal{S} are the action and state spaces, respectively, and the weights are θ . An action is sampled using $a_t \sim \pi(\cdot | s_t; \theta)$. The state transition distribution conditioned on state-action pairs is $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, where probability of the next state s_{t+1} given the current state-action pair s_t, a_t is $p(s_{t+1} | s_t, a_t)$.

Rewards: The reward function, which indicates the value of performing a given action when within a given current state, is defined as $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which, at step t , produces a reward $R_t = r(s_t, a_t)$.

Trajectories: Following the state transition distribution for sampling next states, the policy for sampling actions and the reward function for sampling rewards, we can collect a trajectory of state-action-reward triplets over multiple time-steps $\tau = (s_1, a_1, R_1, \dots, s_T, a_T, R_T)$.

Policy optimisation: The policy $\pi(\cdot; \theta)$, modelled as a neural network, with parameters θ , predicts parameters of a categorical distribution from which an action is to be sampled $a_t \sim \pi(\cdot | s_t; \theta)$. For a policy followed to collect a trajectory, return is given by $R(\tau) = \sum_{k=0}^T \gamma^k R_{t+k}$, which indicates return over trajectory τ . γ is a discount factor for future rewards. The central optimisation problem to find optimal policy parameters θ^* then becomes:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{\tau \sim \pi(\cdot; \theta)} [R(\tau)] \quad (1)$$

Expectation over trajectories is computed using multiple trajectories from agent-environment interactions, where $\tau \sim \pi(\cdot; \theta)$ denotes trajectory sampling using policy π .

Multi-agent self-play environments: In self-play, if observed states for the two competing agents are of the same form, from the point-of-view of the competitors, then sampling an action for an observed state for the competitor may also be achieved using $a_t \sim \pi(\cdot | s_t; \theta)$.

At this point, we introduce subscripts t_a and t_b where a and b indicate the agent a or the competitor agent b , respectively. In this case, the states, actions and rewards for agent a may be denoted as $s_{t_a}, a_{t_a}, R_{t_a}$ and those for agent b as $s_{t_b}, a_{t_b}, R_{t_b}$. Trajectories for agents a and b are denoted as τ_a and τ_b . The optimisation problem is then modified:

$$\theta^* = \arg \max_{\theta} \left\{ \mathbb{E}_{\tau_a \sim \pi(\cdot; \theta)} [R(\tau_a)] + \mathbb{E}_{\tau_b \sim \pi(\cdot; \theta)} [R(\tau_b)] \right\} \quad (2)$$

After each termination, two trajectories τ_a and τ_b are collected, one from each side. Using experience from both sides leads to a diverse range of experiences (from both the winning and losing side), potentially aiding generalisability.

Competitor sampling: While, the section above describes vanilla self-play, fictitious self-play may use a historic version of the policy denoted $\pi_h(\cdot; \theta_h)$, where h denotes the training iteration from which the historic policy is sampled, when the current iteration is denoted as H , i.e., $h \in \{0, 1, 2, \dots, H\}$, where H is incremented by one on each training iteration. So agent a is always denoted as $\pi_H(\cdot; \theta_H)$, using the most up-to-date policy or best response.

For fictitious self-play, the competitor agent b is denoted as $\pi_h(\cdot; \theta_h)$, where $h \sim \rho(h)$ and $\rho(h) = 1/H$ specifying a uniform distribution between 0 to H , over all historic policies. That is to say that a competitor is chosen randomly from all past versions of the policy.

For prioritised fictitious self-play, we propose to use a truncated normal distribution for competitor sampling from historic policies. The competitor agent b is denoted as $\pi_h(\cdot; \theta_h)$, where $h \sim \rho(h)$, where $\rho(h)$ takes the form:

$$\rho(h) = \frac{4}{H} \cdot \frac{\psi\left(\frac{4h-2H}{H}\right)}{\Phi(-1) - \Phi(-2)} \quad (3)$$

Where $\psi(v) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}v^2}$; and $\Phi(v) = \frac{1}{2}(1 + \text{erf}(v/\sqrt{2}))$; and $\text{erf}(v)$ is the Gauss error function.

This is our proposed prioritised fictitious self-play. The distribution that we use is more commonly referred to as a truncated Gaussian distribution, where we use $H/2$ as the mean, $H/4$ as the standard deviation and 0 and H as the bounds. This strategy is likely to select opponents that are weaker than the agent but still allows much weaker or the best response to be sampled for a wide range of experiences to be collected. At the same time, a large amount of collection time is not dedicated to much weaker opponents as in fictitious self-play and cyclic learning / local optima, encountered in vanilla self-play, are avoided as collection time for nearly equally skilled opponents is also reduced.

Agent optimisation: Optimisation for both fictitious or prioritised fictitious self-play are the same as vanilla self-play, with the only difference being that for vanilla self-play, for trajectory τ_a and τ_b collection, the actions for both agent a and agent b are sampled using $a_{t_a} \sim \pi_H(\cdot|s_{t_a};\theta_H)$ and $a_{t_b} \sim \pi_H(\cdot|s_{t_b};\theta_H)$, respectively (i.e., using the most recent or best-response policy), with states for each still being sampled using the state transition distribution p and rewards for each also still computed using the reward function r . For fictitious or prioritised fictitious self-play, actions for agent a and agent b are sampled using $a_{t_a} \sim \pi_H(\cdot|s_{t_a};\theta_H)$ and $a_{t_b} \sim \pi_h(\cdot|s_{t_b};\theta_h)$, respectively (i.e., using the most recent or best-response policy for agent a and sampling a competitor agent b from historic policies), where $h \sim \rho(h)$, and $\rho(h)$ is defined based on whether self-play is fictitious when a uniform distribution over historic policies is used, or prioritised fictitious when a non-uniform distribution over historic policies is used.

B. Gamified segmentation using self-play

An overview of gamified segmentation: Gamified segmentation consists of two agents that select patches from within an image, such that all ROI-containing patches are selected, where the object presence detector acts as a mechanism to score the patch selection by opposing agents. The score is based on the raw classification probability of the patch, from the object presence detector, which serves as a proxy for likelihood of ROI presence within the patch. The game is terminated by means of a termination signal delivered by one of the agents, when it deems all ROI-containing patches to be exhausted. Rewards are given for each move, where the winning patch at each step is determined by the object presence detector, as well as at termination to determine the winning side based on if the opponent can present a patch with a higher score or not, as determined by the object presence detector.

Object presence detector: The object presence detector $f(\cdot; w) : \mathcal{X} \rightarrow [0, 1]$ predicts object or ROI presence (zero for absence and one for presence) within an image or image

Data: Image-label pairs) $\{x_i, y_i\}_{i=1}^N$
Result: Trained RL policy $\pi(\cdot; \theta^*)$

Train object presence detector using image-label pairs (eq. 4) to get $f(\cdot; w^*)$;
 Set current policy number $H = 0$

while not converged do

Randomly sample an image x_i ;
 Start at $t = 0$;
 Divide image x_i into patches $\{x_{i,p}\}_{p=1}^P$;
 Set agent a to have policy $\pi_H(\cdot; \theta_H)$;
 Sample policy $\pi_h(\cdot; \theta_h)$ for agent b (eq. 3);
 Set the initial state as $s_0 = \{x_{i,p}\}_{p=1}^P$;
 Sample action $a_{0_a} \sim \pi_H(\cdot|s_0)$ for agent a ;
 Select patch $x_{i,p_a,0}$ and erase to update s_0 ;
 Sample action $a_{0_b} \sim \pi_h(\cdot|s_0)$ for agent b ;
 Select patch $x_{i,p_b,0}$ and erase to update s_0 ;
 Compute reward R_{0_a} for agent a (eq. 5);
 Compute reward R_{0_b} for agent b (eq. 5 with inverted inequalities);

for $t \leftarrow 1$ to t_{max} do

Carry $s_t = s_{t-1}$ from previous iteration ;
 Sample action $a_{t_a} \sim \pi_H(\cdot|s_t)$ for agent a ;
 Select patch $x_{i,p_a,t}$ and erase to update s_t ;
 Sample action $a_{t_b} \sim \pi_h(\cdot|s_t)$ for agent b ;
 Select patch $x_{i,p_b,t}$ and erase to update s_t ;
 Compute reward R_{t_a} for agent a (eq. 5);
 Compute reward R_{t_b} for agent b (eq. 5 with inverted inequalities);
 Break iterating t if $a_{t,term} = 1$; set $t = t_{end}$;

end

Once $R_{t_a=1:t_{end}}$ and $R_{t_b=1:t_{end}}$ collected, update policy using gradient ascent (eq. 2);
 Update current policy number $H = H + 1$

end

Algorithm 1: RL self-play for WSSS.

patch $x \in \mathcal{X}$, where \mathcal{X} is the domain of images and w are the weights. The only training data used to train f are image-label pairs $\{x_i, y_i\}_{i=1}^N$. Binary cross-entropy loss is used for training: $L_f(y_i, f(x_i; w)) = (y_i) \log(f(x_i; w)) - (1 - y_i) \log(1 - f(x_i; w))$. And the optimisation problem is to find optimal weights w^* , by means of gradient descent:

$$w^* = \arg \min_w \mathbb{E} [L_f(y, f(x; w))] \quad (4)$$

After training, the classifier predicts raw classification probability for the binary object presence classification task $f(x_i; w^*) = [0, 1]$ for image x_i , which is a proxy for likelihood of object presence. Note that while training, only images and image-level classification labels of object presence are used, however, inference may be on image patches that can be interpolated to the image size used for training, regardless of patch size. This enables quantification of the likelihood of object presence within an image patch, which we will use as a reward generator to train our RL self-play segmentation

framework, as outlined in the following sections. This kind of task-based reward for RL is inspired by previous works [13], [70]–[73], [102], [110].

The environment: The gamified segmentation environment encapsulates the image to-be-segmented $x_i \in \mathcal{X}$ and the object presence detector $f(\cdot; w^*)$ with pre-trained weights w^* . All states, actions and rewards are defined from the perspective of agent a , however, definitions may be adapted to agent b e.g., by inverting inequalities defined in rewards. Subscripts a are omitted from some terms for readability, however, are added where ambiguous or where both agents a and b are involved. A single time-step t encompasses moves from both sides, i.e., one patch selection action each (see Algo. 1).

States: The observed state for the environment is an image x_i divided into P patches using a grid of pre-defined size, $s_t = \{x_{i,p}\}_{p=1}^P$, where $x_{i,p}$ is a patch of image x_i .

Actions: We use a P -dimensional discrete action space for patch selection $a_{t,\text{patch}} \in \{1, 2, \dots, P\}$, which selects one of P patches, and a 2-dimensional discrete action space for the termination action $a_{t,\text{term}} \in \{0, 1\}$ (one for termination and zero otherwise). The final action space is then given by $a_t = [a_{t,\text{patch}}, a_{t,\text{term}}]^\top \in \mathcal{A}$. The patch number selected by agent a is defined as p_a and for agent b as p_b , otherwise a subscript of a or b at the action, state or reward indicates the respective agent that is responsible.

Patch selection At any time-step t , a patch selected by agent a is denoted as $x_{i,p_a,t}$, which once selected, is replaced by a zero-intensity patch (pixel values equal to zero) to update observed state s_t , i.e., $x_{i,p_a,t} = \mathbf{O}$. This is similar to commonly-used erasure-based WSSS works [12], [30], [31], [41], [45], [101], such that selecting that patch again is not advantageous in terms of score, since the object presence detector is likely to produce a low score for any zero-intensity patch, in addition to an explicit reward dis-encouraging such selections (see below). The observed state is modified twice at each time-step, once by each of the two agents. After this erasure of selected patches, the updated state is used as the next state s_{t+1} .

Rewards: Reward at each time-step has four parts:

1) the patch reward:

$$R_{t,\text{patch}} = \begin{cases} +1 & \text{if } f(x_{i,p_a,t}; w^*) \geq f(x_{i,p_b,t}; w^*) \\ -1 & \text{otherwise} \end{cases}$$

Which indicates the winning patch at each time-step, as scored by the object presence detector, where the higher raw classification probability for a patch wins the turn. Note that the game is sequential turn-based where one agent plays before another to prevent selection of the same patches by the two agents. This makes the setup a zero-sum game.

2) the terminal reward:

$$R_{t,\text{term}} = \begin{cases} 0 & \text{if } a_{t,\text{term}} = 0 \\ \text{otherwise:} \\ +1 & \text{if } f(x_{i,p_a,t}; w^*) \geq f(x_{i,p_b,t}; w^*) \\ -1 & \text{if } f(x_{i,p_a,t}; w^*) < f(x_{i,p_b,t}; w^*) \end{cases}$$

Which decides the final game winner at termination such that if a termination signal is delivered by agent a and agent a loses the turn, a negative reward is given otherwise if agent a wins the terminal time-step, a positive reward is delivered. The

terminal reward plays an important role in preventing over- and under- segmentation otherwise encountered with image-level supervision (see Sec. II-B and I).

3) the repetition reward:

$$R_{t,\text{rep}} = \begin{cases} -1 & \text{if } x_{i,p_a,t} = \mathbf{O} \\ 0 & \text{otherwise} \end{cases}$$

Which discourages selection of an already selected patch as these are replaced by zero-intensity patches \mathbf{O} in state s_t .

4) the iteration bounding reward:

$$R_{t,\text{iter}} = \begin{cases} -1 & \text{if } t_{\min} < t < t_{\max} \text{ and } a_{t,\text{term}} = 1 \\ 0 & \text{otherwise} \end{cases}$$

Which prevents early termination before t_{\min} and also late termination after t_{\max} .

In practice $R_{t,\text{rep}}$ and $R_{t,\text{iter}}$ act as shaping to make learning faster, however, are not necessary (see Table IV).

The final reward for the time-step t is thus formed using all four components in a weighted manner:

$$R_t = R_{t,\text{patch}} + 100R_{t,\text{term}} + R_{t,\text{rep}} + R_{t,\text{iter}} \quad (5)$$

Where the weights were tuned on a small subset of the PASCAL VOC 2012 train set (see Sec. IV-A).

Agent-environment interactions: The above-described states, actions and rewards are used to train a policy (eq. 2). An overview is outlined in Algo. 1 and Fig. 2. At inference, we run the policy without opposition, or with a fixed opponent policy which always selects an artificially-placed zero-intensity patch. Inference to obtain a final segmentation map is done via the agent selecting patches containing the object without opposition. The grid of patches can be shifted by one pixel, or more, in order to run inference again to conduct a majority vote over pixel selections to get a pixel-level segmentation map.

IV. EXPERIMENTS AND RESULTS

A. Datasets for evaluation

PASCAL VOC 2012 (VOC) [16] has 20 classes with 1,464, 1,449 and 1,456 images in the train, validation and test sets. As common practice [11], [32], [111], the train set was augmented with 10,582 images [26], with 1024 used for monitoring performance, and the validation and test sets were held out and used to report results (test results obtained from evaluation server).

MS COCO 2014 (COCO) [58] has 80 classes with 80k and 40k images in the train and validation sets (20k images used from train set for monitoring performance). Validation set was held out and used to report results, as common practice [11], [111].

Male Pelvic Structures (CMPS) [54] has 8 classes with 471 and 118 samples in the train and test sets (64 samples from train set used for monitoring performance). Test set was held out and used to report performance. We use 2D slices with each 3D sample containing approximately 20–40 slices.

Liver Tumour (LiTS) [6] has 1 class (tumour) with 107 and 24 samples in the train and test sets (24 images from train set used for monitoring performance). The test set was held out and used to report performance. We use 2D slices with each 3D sample containing approximately 200–400 slices.

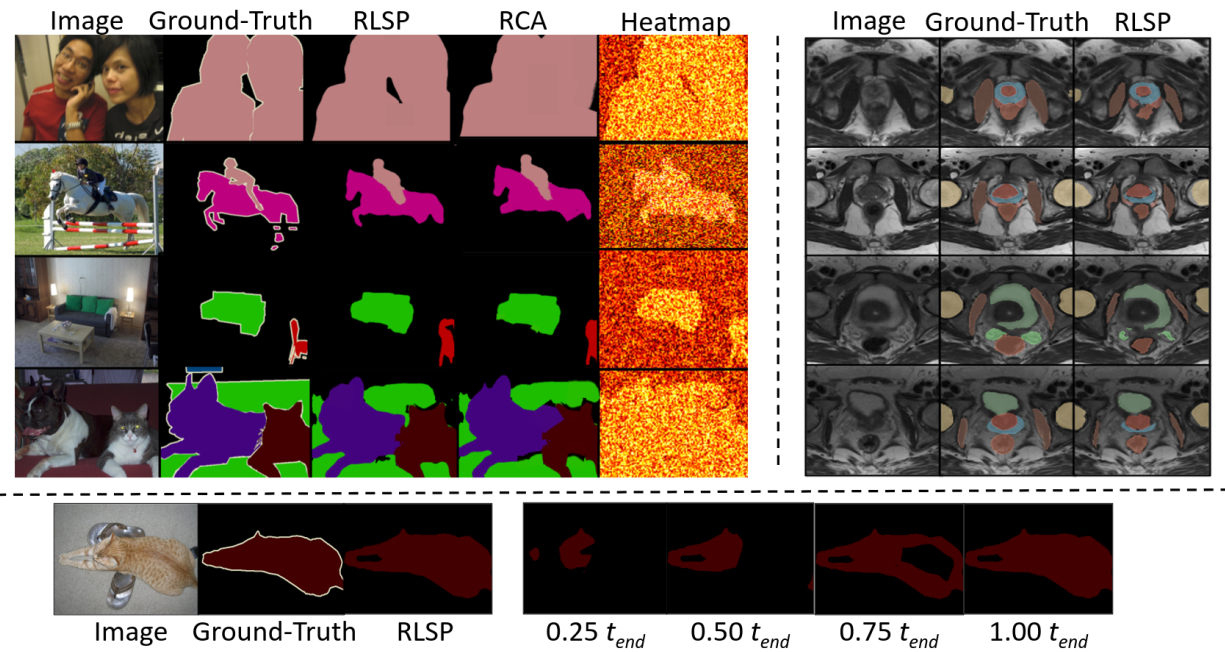


Fig. 3. Left: VOC; Right: CMPS. ‘Heatmap’: object presence detector score heatmap. Bottom: VOC; time-points at inference. RLSP uses pixel-level majority vote by running inference thrice.

1) *Monitoring performance*: Performance can be monitored during RL training by examining the obtained rewards from agent a . We used trials of 32 episodes (each episode being t_a from 1 to t_{end}) to train and 8 episodes to compute rewards for performance monitoring. This trial-based training and reward computation was performed until convergence. We observed an increase in the performance-monitoring-rewards until a plateau and set a convergence criterion to stop training (e.g., minimum-delta of $1e - 5$ across a moving average with window size of 16 trials, for VOC). Additionally, performance can be monitored while training the object presence detector by setting a convergence criterion on the performance metric (e.g., minimum-delta of $1e - 6$ across 4 consecutive epochs on mIoU on the performance-monitoring-samples, for VOC).

B. Network architectures and hyper-parameters

We use ResNet38 [27], [95] as a backbone for our object presence detector and agent. We train a separate network for identifying object presence or absence, for each class (threshold 0.5 used at inference, to decide classes present in the image). Hyper-parameter values are summarised in Table I.

The object presence detector is trained as a binary classification network (of object presence) for each class that requires segmentation, when there are more than two classes in a dataset. The training data is split into positive and negative images where the positive samples are the ROI-containing samples belonging to the class-of-interest and the negative samples are all other samples in the dataset. For efficient training of the object presence detector for each class, 20% data from each of the remaining classes were randomly sampled and used as negative samples for training. This was done for large datasets including VOC and COCO, to reduce

overall training time and maintain a class-balance between positive and negative samples [52]. Each class that is to-be-segmented requires a separate object presence detector. This training mechanism ensures that the object presence detector, for each class, can learn to identify object presence vs absence, which serves as a reward signal for our proposed self-play segmentation approach. At inference, we use the object presence detector to get binary object presence decisions per-image, for every class. Subsequently, we only run self-play with an object presence detector from a specific class if the class probability for the sample is greater than 0.5, otherwise we assume no object of the specific class is present within an image.

We also refer the readers to the open-source repository for further implementation details.

TABLE I
HYPER-PARAMETER VALUES.

Hyper-parameter	Value
Object presence detector	
Input-size	256×256
Architecture	ResNet38
Batch-size	64
Optimiser	Adam
Self-play	
RL algorithm	Policy Gradient
Input-size	256×256
Architecture	ResNet38
Batch-size	128
Optimiser	Adam
Gamma γ	0.96
t_{min}	4
t_{max}	1024
Patch-size	6×6

TABLE II
SEGMENTATION RESULTS. VOC, COCO (mIoU); CMPS, LiTS (mDSC).

Method	Backbone	VOC (val)	VOC (test)	COCO	CMPS (RE)	LiTS (RE)
SEC [40]	VGG16	-	-	22.4	-	-
IRN [1]	ResNet50	63.5	64.8	32.6	-	-
IAL [89]	VGG16	-	-	27.7	-	-
ICD [17]	ResNet101	64.1	64.3	-	-	-
SCE [8]	ResNet101	66.1	65.9	-	-	-
SEAM [92]	ResNet38	64.5	65.7	31.9	-	-
BES [10]	ResNet101	65.7	66.6	-	-	-
MCIS [83]	ResNet101	66.2	66.9	-	-	-
CONTA [108]	ResNet101	66.1	66.7	33.4	-	-
LIID [61]	ResNet101	66.5	67.5	-	-	-
A2GNN [105]	ResNet101	66.8	67.4	-	-	-
AdvCAM [46]	ResNet101	68.1	68.0	-	-	-
CDA [82]	ResNet38	66.1	66.8	-	-	-
ECS [84]	ResNet38	66.6	67.6	-	-	-
CSE [42]	ResNet38	68.4	68.2	36.4	-	-
CPN [109]	ResNet38	67.8	68.5	-	-	-
USAGE [65]	ResNet38	71.9	72.8	44.3	-	-
CDL [107]	ResNet101	72.4	72.9	48.7	-	-
MIL (RE) [50]	VGG16	67.9	-	37.9	52.1	67.4
RCA [111]	ResNet38	72.2	72.8	36.8	53.9	68.8
SIPE [11]	ResNet38	68.2	69.5	43.6	53.7	66.1
MCT [98]	ResNet38	71.9	71.6	42.0	54.1	67.1
ACR [43]	ResNet38	72.4	71.9	45.3	53.7	67.9
MARS [36]	ResNet101	77.7	77.2	49.4	54.2	68.6
RLSP (ours)	ResNet38	78.9	78.7	49.9	58.8	70.7

C. Comparisons with the state-of-the-art

We compare our proposed RL self-play (RLSP) against multiple proposed CAM or MIL-based WSSS works that utilise only image-level labels for training (citations provided where appropriate). Re-implemented methods are denoted with a suffix ‘RE’.

Previous works do not report measures of spread so statistical tests were not conducted. For re-implemented methods, we report t-test results (significance $\alpha = 0.05$).

We report the commonly used overlap measures mean intersection over union (mIoU), for VOC and COCO, or mean binary Dice score (mDSC), for CMPS and LiTS, dependant on what compared methods have reported.

Table II shows that our method outperformed recent state-of-the-art (SOTA) WSSS methods that only use image-level classification labels for training, on all tested datasets. We observe improvements of 1.2 mIoU points on VOC (validation set), 1.7 mIoU points on VOC (test set), 0.5 mIoU points on COCO, 4.6 mDSC points on the CMPS, and 1.9 mDSC points on the LiTS dataset. For the CMPS and LiTS datasets, differences are statistically significant (p-values < 0.001). Modifying the backbone of our method for VOC did not substantially impact performance i.e., changing from ResNet38 to ResNet101 or VGG16 marginally reduced performance, by 0.5 or 1.2 mIoU points, respectively.

We observe that our proposed method performs best under scenarios where the ROI is smaller than (less than half) the whole FOV of the image (qualitative samples in Fig. 3 and 4).

Training and inference times varied with the number of classes and approximate ROI size in the dataset e.g., for VOC, we used approximately 450 GPU hours for training on an Nvidia Tesla V100 GPU (48 hours pre-training object presence detector), and 2.5 seconds for inference per image.

Performance improvements across a wide variety of datasets demonstrate the applicability of our framework.

D. Ablation studies

Table III shows impact of increasing patch-size. This leads to improved performance, until a plateau, which is intuitive as our classifier tends to perform well when input FOV matches the training FOV i.e., as the input patches get larger and closer to whole images. Future development could explore an adaptive object presence detector that is trained together with the agents for optimal performance for a given patch size. However, since VOC had a wide variety of FOVs in the training data, impact of patch size was limited.

TABLE III
RESULTS FOR VARYING PATCH SIZE. VOC (VAL) (mIoU).

Patch size	VOC
2×2	76.4
4×4	76.6
6×6	78.9
8×8	78.3
10×10	78.4
12×12	78.8
16×16	78.7

Table IV, shows impact of removal of components from our proposed framework. Non-self-play variants are trained with a reward based on a threshold 0.8 in place of agent b , i.e., setting $f(x_{i,p_b,t}; w^*) = 0.8$. Omission of shaping refers to repetition and/ or iteration bounding rewards i.e., $R_{t,rep} = 0$ and/ or $R_{t,iter} = 0$. We observe significant negative impact without self-play, however, impact is limited when shaping rewards are omitted. Note that reward shaping reduced training times by approximately 8 hours.

TABLE IV
RESULTS FOR ABLATION STUDY. VOC (VAL) (mIoU).

SP	SR		VOC
	$R_{t,rep}$	$R_{t,iter}$	
✗	✗	✗	68.8
✗	✓	✓	70.3
✓	✗	✗	75.7
✓	✓	✗	76.4
✓	✗	✓	77.7
✓	✓	✓	78.9

E. Comparisons to baselines

Table V, shows comparisons to various baseline models such as a fully-supervised network (trained with pixel-level ground truth), sliding window approach (using the object presence detector with a commonly used threshold of 0.5 for segmentation), and a single-agent formulation denoted as non-SP, with varying thresholds used for selecting patches. We demonstrate superior performance of our proposed framework compared to all tested variants trained with weak image-level labels. Furthermore, we report results for different seeds and show that our training scheme is robust to seed selection with a variance (standard deviation) of 0.47 over 8 random seeds. Given that patch-size has limited impact and that compared to sliding window approach without any RL, we obtain a higher performance for RLSP, combined with results from the ablation study, we conclude that self-play is the largest positive contributor towards performance in our proposed framework.

TABLE V
COMPARISONS TO BASELINES. VOC (VAL) (mIoU).

Variant	VOC
Fully-supervised	81.2
Non-SP ($f(x_{i,p_b,t}; w^*) = 0.8$)	70.3
Non-SP ($f(x_{i,p_b,t}; w^*) = 0.5$)	69.8
Non-SP ($f(x_{i,p_b,t}; w^*) = 0.2$)	68.5
Sliding-window	64.9
RLSP (seed 1)	78.9
RLSP (seed 2)	77.9
RLSP (seed 3)	78.5

Performance improvements in overlap measures (mIoU and mDSC) compared to other methods (see Table II), or variants (see Table V) with sliding windows or omitted self-play, indicate the ability of our proposed method to provide effective segmentation. These overlap measures are reduced by over- or under-segmentation and thus we conclude that our method, with the highest overlap measure amongst tested WSSS methods, learns an effective termination strategy which minimises over- or under-segmentation compared to other methods and tested variants.

F. Quantifying over- and under-segmentation

Additionally, in Table VI we quantify over- and under-segmentation by reporting the proportion of true-positive (TP), true-negative (TN), false-positive (FP) and false-negative (FN) pixel values, averaged over the set (normalised to a sum of 100). The higher TP and TN and lower FP and FN for our method further indicate that over- and under-segmentation were minimised compared to other tested methods.

TABLE VI
COMPARING OVER-/ UNDER-SEGMENTATION. VOC (VAL).

Method	mIoU	TP	TN	FP	FN
Sliding-window	64.9	63	15	10	12
Non-SP	70.3	65	15	10	12
MARS [36]	77.7	67	17	8	8
RLSP	78.9	69	18	6	7

G. Impact of network architectures and RL training algorithms

Table VII shows variants of RLSP with different network architectures used as the backbones for the object presence detector and RL agents and different RL algorithms used for training.

TABLE VII
IMPACT OF NETWORK ARCHITECTURES AND RL ALGORITHMS USED FOR TRAINING SHOWING ACCURACY OF THE OBJECT PRESENCE DETECTOR AND SEGMENTATION PERFORMANCE (mIoU) FOR RLSP.

Detector	Self-play	RL algorithm	Accuracy	VOC (mIoU)
Testing different backbones				
ResNet38	ResNet38	PG	89.7	78.9
ResNet101	ResNet101	PG	87.5	78.4
VGG16	VGG16	PG	86.4	77.7
AlexNet	AlexNet	PG	80.9	72.3
Testing different RL algorithms				
ResNet38	ResNet38	PG	89.7	78.9
ResNet38	ResNet38	PPO	89.7	78.7
ResNet38	ResNet38	DDPG	89.7	77.8
Testing different backbone combinations				
ResNet38	ResNet38	PG	89.7	78.9
AlexNet	ResNet38	PG	80.9	73.5
ResNet101	ResNet38	PG	87.5	78.8
ResNet38	AlexNet	PG	89.7	77.5
ResNet38	ResNet101	PG	89.7	78.6

The results in Table VII show that our method is not only robust to various backbone networks but also to different RL algorithms, which indicates that the developed reward mechanisms provide a consistent robust supervision signal to train a WSSS system.

H. Using pseudo ground truth for segmentation network training

In addition to the predictions from RLSP used for evaluation so far, i.e., obtaining segmentation directly as an output from the proposed self-play mechanism, it is also possible to train a segmentation network using the output of RLSP as pseudo ground truth labels, similar to previous work [20]. This can allow for inference time to be reduced while maintaining the weakly-supervised nature of the proposed approach, as the segmentation labels used to train the segmentation network are derived from the WSSS RLSP as opposed to being manually defined by a human observer. We use 50% of the validation set pseudo labels to train the segmentation network and the remaining 50% to report performance, as summarised in Table VIII.

These results, from Table VIII, show that training a segmentation network using pseudo ground truth labels from RLSP

TABLE VIII
RESULTS FOR USING PSEUDO GROUND TRUTH LABELS FROM RLSP TO TRAIN A SEGMENTATION NETWORK. VOC (VAL) (mIoU).

Segmentation Net	Inference Time / s	VOC (val)
None (simple RLSP)	2.500	78.5
ConvNet [62]	0.120	75.2
UNet [68]	0.150	77.8
DeepLab v3 [9]	0.180	78.0



Fig. 4. VOC test set segmentation results (left to right: image; ground truth; proposed).

has the potential to reduce inference time while retaining a majority of the performance benefit compared with other WSSS methods. Although the inference time for RLSP itself is higher than previous methods, the comparative performance improvements and ability to reduce inference time via segmentation network training using RLSP as pseudo-labels opens up real-time applications at deployment.

V. DISCUSSION AND CONCLUSION

The results presented in Sec. IV demonstrate the effectiveness of our proposed method, trained only using weak image-level classification labels, by showing improved performance compared to competing methods, ablated variants and other

baselines. The tasks on which our framework is evaluated demonstrate a wide range of applications including two multi-class segmentation computer-vision datasets and two medical imaging datasets, showing applicability to various types of images and ROIs.

Direct comparisons to other methods, in terms correct classifications of pixels containing ROIs, using true-positive, true-negative, false-positive and false-negative ratios, also demonstrate the ability of the method to reduce over- or under-segmentation as it shows higher true-negative and -positive rates and lower false-negative and -positive rates compared to other tested methods.

The task-based rewards, in the self-play set-up, proposed in this work, open up potential avenues for future work where weak supervision signals may be derived for tasks such as image-alignment, object localisation in 3D space, or anomaly detection. Moreover, extending our task-based rewards holds promise for facilitating meta-learning within a self-play context. In such set-ups two task networks may engage in competitive interactions supervised by an outer-level network, all trained simultaneously, enhancing the modelling of intricate interdependencies across these functions. In conclusion, our novel self-play framework opens up further avenues for research into supervisory signals and meta-learning configurations.

In this work, we propose a gamified WSSS such that competing agents segment a ROI from an image trying to minimise over- or under-segmentation. The weak supervision signals used during training are derived only from an object presence detector, which is able to classify object presence within a selected image patch, and itself is trained only using image-level binary classification labels of object presence. Extensive experiments on four datasets show that our proposed framework outperforms recently proposed SOTA methods on two well-known WSSS benchmarks as well as on two real-world medical datasets.

ACKNOWLEDGEMENTS

This work is supported by the EPSRC grant [EP/T029404/1]; Wellcome/EPSC Centre for Interventional and Surgical Sciences [203145Z/16/Z]; the EPSRC CDT in i4health [EP/S021930/1]; and the International Alliance for Cancer Early Detection, an alliance between Cancer Research UK [C28070/A30912; C73666/A31378], Canary Center at Stanford University, the University of Cambridge, OHSU Knight Cancer Institute, University College London and the University of Manchester. For the purpose of open access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

REFERENCES

- [1] Ahn, J., Cho, S., Kwak, S.: Weakly supervised learning of instance segmentation with inter-pixel relations. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2209–2218 (2019)
- [2] Araslanov, N., Roth, S.: Single-stage semantic segmentation from image labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4253–4262 (2020)

- [3] Bai, Y., Jin, C.: Provable self-play algorithms for competitive reinforcement learning. In: International conference on machine learning. pp. 551–560. PMLR (2020)
- [4] Bai, Y., Jin, C., Yu, T.: Near-optimal reinforcement learning with self-play. *Advances in neural information processing systems* **33**, 2159–2170 (2020)
- [5] Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., Mordatch, I.: Emergent tool use from multi-agent autocurricula. *arXiv preprint arXiv:1909.07528* (2019)
- [6] Bilic, P., Christ, P., Li, H.B., Vorontsov, E., Ben-Cohen, A., Kaissis, G., Szeskin, A., Jacobs, C., Mamani, G.E.H., Chartrand, G., et al.: The liver tumor segmentation benchmark (lits). *Medical Image Analysis* **84**, 102680 (2023)
- [7] Brown, G.W.: Iterative solution of games by fictitious play. *Act. Anal. Prod Allocation* **13**(1), 374 (1951)
- [8] Chang, Y.T., Wang, Q., Hung, W.C., Piramuthu, R., Tsai, Y.H., Yang, M.H.: Weakly-supervised semantic segmentation via sub-category exploration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8991–9000 (2020)
- [9] Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587* (2017)
- [10] Chen, L., Wu, W., Fu, C., Han, X., Zhang, Y.: Weakly supervised semantic segmentation with boundary exploration. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16. pp. 347–362. Springer (2020)
- [11] Chen, Q., Yang, L., Lai, J.H., Xie, X.: Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4288–4298 (2022)
- [12] Choe, J., Lee, S., Shim, H.: Attention-based dropout layer for weakly supervised single object localization and semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence* **43**(12), 4256–4271 (2020)
- [13] Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501* (2018)
- [14] Dai, J., He, K., Sun, J.: Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: Proceedings of the IEEE international conference on computer vision. pp. 1635–1643 (2015)
- [15] DiGiovanni, A., Zell, E.C.: Survey of self-play in reinforcement learning. *arXiv preprint arXiv:2107.02850* (2021)
- [16] Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. *International journal of computer vision* **111**, 98–136 (2015)
- [17] Fan, J., Zhang, Z., Song, C., Tan, T.: Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4283–4292 (2020)
- [18] Fan, J., Zhang, Z., Tan, T., Song, C., Xiao, J.: Cian: Cross-image affinity net for weakly supervised semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 10762–10769 (2020)
- [19] Feng, D., Haase-Schütz, C., Rosenbaum, L., Hertlein, H., Glaeser, C., Timm, F., Wiesbeck, W., Dietmayer, K.: Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems* **22**(3), 1341–1360 (2020)
- [20] Feng, J., Wang, X., Liu, W.: Deep graph cut network for weakly-supervised semantic segmentation. *Science China Information Sciences* **64**, 1–12 (2021)
- [21] Foulds, J., Frank, E.: A review of multi-instance learning assumptions. *The knowledge engineering review* **25**(1), 1–25 (2010)
- [22] Ge, W., Guo, S., Huang, W., Scott, M.R.: Label-penet: Sequential label propagation and enhancement networks for weakly supervised instance segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3345–3354 (2019)
- [23] Haarnoja, T., Moran, B., Lever, G., Huang, S.H., Tirumala, D., Wulfmeier, M., Humplik, J., Tunyasuvunakool, S., Siegel, N.Y., Hafner, R., et al.: Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *arXiv preprint arXiv:2304.13653* (2023)
- [24] Hammar, K., Stadler, R.: Finding effective security strategies through reinforcement learning and self-play. In: 2020 16th International Conference on Network and Service Management (CNSM). pp. 1–9. IEEE (2020)
- [25] Han, C., Lin, J., Mai, J., Wang, Y., Zhang, Q., Zhao, B., Chen, X., Pan, X., Shi, Z., Xu, Z., et al.: Multi-layer pseudo-supervision for histopathology tissue semantic segmentation using patch-level classification labels. *Medical Image Analysis* **80**, 102487 (2022)
- [26] Hariharan, B., Arbeláez, P., Bourdev, L., Maji, S., Malik, J.: Semantic contours from inverse detectors. In: 2011 international conference on computer vision. pp. 991–998. IEEE (2011)
- [27] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- [28] Heinrich, J., Lanctot, M., Silver, D.: Fictitious self-play in extensive-form games. In: International conference on machine learning. pp. 805–813. PMLR (2015)
- [29] Heinrich, J., Silver, D.: Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121* (2016)
- [30] Hou, Q., Jiang, P., Wei, Y., Cheng, M.M.: Self-erasing network for integral object attention. *Advances in Neural Information Processing Systems* **31** (2018)
- [31] Huang, Y., Yang, X., Zou, Y., Chen, C., Wang, J., Dou, H., Ravikummar, N., Frangi, A.F., Zhou, J., Ni, D.: Flip learning: Erase to segment. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24. pp. 493–502. Springer (2021)
- [32] Huang, Z., Wang, X., Wang, J., Liu, W., Wang, J.: Weakly-supervised semantic segmentation network with deep seeded region growing. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7014–7023 (2018)
- [33] Hwang, J., Kim, S., Son, J., Han, B.: Weakly supervised instance segmentation by deep community learning. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 1020–1029 (2021)
- [34] Jia, Z., Huang, X., Eric, I., Chang, C., Xu, Y.: Constrained deep weak supervision for histopathology image segmentation. *IEEE transactions on medical imaging* **36**(11), 2376–2388 (2017)
- [35] Jiang, P.T., Hou, Q., Cao, Y., Cheng, M.M., Wei, Y., Xiong, H.K.: Integral object mining via online attention accumulation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2070–2079 (2019)
- [36] Jo, S., Yu, I.J., Kim, K.: Mars: Model-agnostic biased object removal without additional supervision for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2304.09913* (2023)
- [37] Kajiura, N., Kosugi, S., Wang, X., Yamasaki, T.: Self-play reinforcement learning for fast image retargeting. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 1755–1763 (2020)
- [38] Kim, B., Han, S., Kim, J.: Discriminative region suppression for weakly-supervised semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1754–1761 (2021)
- [39] Kim, B., Yoo, Y., Rhee, C.E., Kim, J.: Beyond semantic to instance segmentation: Weakly-supervised instance segmentation via semantic knowledge transfer and self-refinement. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4278–4287 (2022)
- [40] Kolesnikov, A., Lampert, C.H.: Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14. pp. 695–711. Springer (2016)
- [41] Kumar Singh, K., Jae Lee, Y.: Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3524–3533 (2017)
- [42] Kweon, H., Yoon, S.H., Kim, H., Park, D., Yoon, K.J.: Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6994–7003 (2021)
- [43] Kweon, H., Yoon, S.H., Yoon, K.J.: Weakly supervised semantic segmentation via adversarial learning of classifier and reconstructor. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11329–11339 (2023)
- [44] Laterre, A., Fu, Y., Jabri, M.K., Cohen, A.S., Kas, D., Hajjar, K., Dahl, T.S., Kerkeni, A., Beguir, K.: Ranked reward: Enabling self-play reinforcement learning for combinatorial optimization. *arXiv preprint arXiv:1807.01672* (2018)

- [45] Lee, J., Kim, E., Lee, S., Lee, J., Yoon, S.: Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5267–5276 (2019)
- [46] Lee, J., Kim, E., Yoon, S.: Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4071–4080 (2021)
- [47] Lee, J., Yi, J., Shin, C., Yoon, S.: Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2643–2652 (2021)
- [48] Lee, S., Lee, M., Lee, J., Shim, H.: Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5495–5505 (2021)
- [49] Leslie, D.S., Collins, E.J.: Generalised weakened fictitious play. *Games and Economic Behavior* **56**(2), 285–298 (2006)
- [50] Li, K., Qian, Z., Han, Y., Eric, I., Chang, C., Wei, B., Lai, M., Liao, J., Fan, Y., Xu, Y.: Weakly supervised histopathology image segmentation with self-attention. *Medical Image Analysis* **86**, 102791 (2023)
- [51] Li, Q., Arnab, A., Torr, P.H.: Weakly- and semi-supervised panoptic segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 102–118 (2018)
- [52] Li, X., Wei, T., Chen, Y.P., Tai, Y.W., Tang, C.K.: Fss-1000: A 1000-class dataset for few-shot segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2869–2878 (2020)
- [53] Li, X., Zhou, T., Li, J., Zhou, Y., Zhang, Z.: Group-wise semantic mining for weakly supervised semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1984–1992 (2021)
- [54] Li, Y., Fu, Y., Gayo, I.J., Yang, Q., Min, Z., Saeed, S.U., Yan, W., Wang, Y., Noble, J.A., Emberton, M., et al.: Prototypical few-shot segmentation for cross-institution male pelvic structures with spatial registration. *Medical Image Analysis* **90**, 102935 (2023)
- [55] Li, Z., Xia, Y.: Deep reinforcement learning for weakly-supervised lymph node segmentation in ct images. *IEEE Journal of Biomedical and Health Informatics* **25**(3), 774–783 (2020)
- [56] Liang, Z., Wang, T., Zhang, X., Sun, J., Shen, J.: Tree energy loss: Towards sparsely annotated semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16907–16916 (2022)
- [57] Lin, D., Dai, J., Jia, J., He, K., Sun, J.: Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3159–3167 (2016)
- [58] Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13. pp. 740–755. Springer (2014)
- [59] Liu, S., Cao, J., Wang, Y., Chen, W., Liu, Y.: Self-play reinforcement learning with comprehensive critic in computer games. *Neurocomputing* **449**, 207–213 (2021)
- [60] Liu, W., Zhang, C., Lin, G., Hung, T.Y., Miao, C.: Weakly supervised segmentation with maximum bipartite graph matching. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2085–2094 (2020)
- [61] Liu, Y., Wu, Y.H., Wen, P., Shi, Y., Qiu, Y., Cheng, M.M.: Leveraging instance-, image-and dataset-level information for weakly supervised instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(3), 1415–1428 (2020)
- [62] Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
- [63] Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., Terzopoulos, D.: Image segmentation using deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(7), 3523–3542 (2021)
- [64] Oh, Y., Kim, B., Ham, B.: Background-aware pooling and noise-aware loss for weakly-supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6913–6922 (2021)
- [65] Peng, Z., Wang, G., Xie, L., Jiang, D., Shen, W., Tian, Q.: Usage: A unified seed area generation paradigm for weakly supervised semantic segmentation. *arXiv preprint arXiv:2303.07806* (2023)
- [66] Portelas, R., Colas, C., Weng, L., Hofmann, K., Oudeyer, P.Y.: Automatic curriculum learning for deep rl: A short survey. *arXiv preprint arXiv:2003.04664* (2020)
- [67] Ramesh, K., Kumar, G.K., Swapna, K., Datta, D., Rajest, S.S.: A review of medical image segmentation algorithms. *EAI Endorsed Transactions on Pervasive Health and Technology* **7**(27), e6–e6 (2021)
- [68] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III* 18. pp. 234–241. Springer (2015)
- [69] Roth, H.R., Yang, D., Xu, Z., Wang, X., Xu, D.: Going to extremes: weakly supervised medical image segmentation. *Machine Learning and Knowledge Extraction* **3**(2), 507–524 (2021)
- [70] Saeed, S.U., Fu, Y., Baum, Z.M., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Barratt, D.C., Hu, Y.: Learning image quality assessment by reinforcing task amenable data selection. In: *International Conference on Information Processing in Medical Imaging*. pp. 755–766. Springer (2021)
- [71] Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Noble, J.A., Barratt, D.C., et al.: Adaptable image quality assessment using meta-reinforcement learning of task amenability. In: *Simplifying Medical Ultrasound: Second International Workshop, ASMUS 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 2*. pp. 191–201. Springer (2021)
- [72] Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Noble, J.A., Barratt, D.C., et al.: Image quality assessment for machine learning tasks using meta-reinforcement learning. *Medical Image Analysis* **78**, 102427 (2022)
- [73] Saeed, S.U., Yan, W., Fu, Y., Giganti, F., Yang, Q., Baum, Z., Rusu, M., Fan, R.E., Sonn, G.A., Emberton, M., et al.: Image quality assessment by overlapping task-specific and task-agnostic measures: application to prostate multiparametric mr images for cancer segmentation. *arXiv preprint arXiv:2202.09798* (2022)
- [74] Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE international conference on computer vision*. pp. 618–626 (2017)
- [75] Shamma, J.S., Arslan, G.: Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria. *IEEE Transactions on Automatic Control* **50**(3), 312–327 (2005)
- [76] Sharma, N., Aggarwal, L.M.: Automated medical image segmentation techniques. *Journal of medical physics/Association of Medical Physicists of India* **35**(1), 3 (2010)
- [77] Shimoda, W., Yanai, K.: Self-supervised difference detection for weakly-supervised semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5208–5217 (2019)
- [78] Shoham, Y., Powers, R., Grenager, T.: Multi-agent reinforcement learning: a critical survey. *Tech. rep., Citeseer* (2003)
- [79] Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al.: Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815* (2017)
- [80] Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., et al.: A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* **362**(6419), 1140–1144 (2018)
- [81] Song, C., Huang, Y., Ouyang, W., Wang, L.: Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3136–3145 (2019)
- [82] Su, Y., Sun, R., Lin, G., Wu, Q.: Context decoupling augmentation for weakly supervised semantic segmentation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 7004–7014 (2021)
- [83] Sun, G., Wang, W., Dai, J., Van Gool, L.: Mining cross-image semantics for weakly supervised semantic segmentation. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II* 16. pp. 347–365. Springer (2020)
- [84] Sun, K., Shi, H., Zhang, Z., Huang, Y.: Ecs-net: Improving weakly supervised semantic segmentation by using connections between class activation maps. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 7283–7292 (2021)
- [85] Team, O.E.L., Stooke, A., Mahajan, A., Barros, C., Deck, C., Bauer, J., Sygnowski, J., Trebacz, M., Jaderberg, M., Mathieu, M., et al.:

- Open-ended learning leads to generally capable agents. arXiv preprint arXiv:2107.12808 (2021)
- [86] Van Der Ree, M., Wiering, M.: Reinforcement learning in the game of othello: Learning against a fixed opponent and learning from self-play. In: 2013 IEEE symposium on adaptive dynamic programming and reinforcement learning (ADPRL). pp. 108–115. IEEE (2013)
- [87] Vernaza, P., Chandraker, M.: Learning random-walk label propagation for weakly-supervised semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7158–7166 (2017)
- [88] Vinyals, O., Babuschkin, I., Czarnecki, W.M., Mathieu, M., Dudzik, A., Chung, J., Choi, D.H., Powell, R., Ewalds, T., Georgiev, P., et al.: Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature* **575**(7782), 350–354 (2019)
- [89] Wang, X., Liu, S., Ma, H., Yang, M.H.: Weakly-supervised semantic segmentation by iterative affinity learning. *International Journal of Computer Vision* **128**, 1736–1749 (2020)
- [90] Wang, X., You, S., Li, X., Ma, H.: Weakly-supervised semantic segmentation by iteratively mining common object features. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1354–1362 (2018)
- [91] Wang, Y., Tang, H., Huang, L., Pan, L., Yang, L., Yang, H., Mu, F., Yang, M.: Self-play reinforcement learning guides protein engineering. *Nature Machine Intelligence* pp. 1–16 (2023)
- [92] Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X.: Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12275–12284 (2020)
- [93] Wei, Y., Feng, J., Liang, X., Cheng, M.M., Zhao, Y., Yan, S.: Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1568–1576 (2017)
- [94] Wei, Y., Xiao, H., Shi, H., Jie, Z., Feng, J., Huang, T.S.: Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7268–7277 (2018)
- [95] Wu, Z., Shen, C., Van Den Hengel, A.: Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition* **90**, 119–133 (2019)
- [96] Xu, G., Song, Z., Sun, Z., Ku, C., Yang, Z., Liu, C., Wang, S., Ma, J., Xu, W.: Camel: A weakly supervised learning framework for histopathology image segmentation. In: Proceedings of the IEEE/CVF International Conference on computer vision. pp. 10682–10691 (2019)
- [97] Xu, L., Ouyang, W., Bennamoun, M., Boussaid, F., Sohel, F., Xu, D.: Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6984–6993 (2021)
- [98] Xu, L., Ouyang, W., Bennamoun, M., Boussaid, F., Xu, D.: Multi-class token transformer for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4310–4319 (2022)
- [99] Xu, Y., Zhu, J.Y., Eric, I., Chang, C., Lai, M., Tu, Z.: Weakly supervised histopathology cancer image segmentation and classification. *Medical image analysis* **18**(3), 591–604 (2014)
- [100] Yao, Y., Chen, T., Xie, G.S., Zhang, C., Shen, F., Wu, Q., Tang, Z., Zhang, J.: Non-salient region object mining for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2623–2632 (2021)
- [101] Yi, W., Stavrinides, V., Baum, Z., Yang, Q., Barratt, D.C., Clarkson, M.J., Hu, Y., Saeed, S.U.: Boundary-rl: Reinforcement learning for weakly-supervised prostate segmentation in trus images. arXiv preprint arXiv:2308.11376 (2023)
- [102] Yoon, J., Arik, S., Pfister, T.: Data valuation using reinforcement learning. In: International Conference on Machine Learning. pp. 10842–10851. PMLR (2020)
- [103] Yun, K., Peng, Y., Samaras, D., Zelinsky, G.J., Berg, T.L.: Exploring the role of gaze behavior and object detection in scene understanding. *Frontiers in psychology* **4**, 917 (2013)
- [104] Zha, D., Xie, J., Ma, W., Zhang, S., Lian, X., Hu, X., Liu, J.: Douzero: Mastering douzihu with self-play deep reinforcement learning. In: international conference on machine learning. pp. 12333–12344. PMLR (2021)
- [105] Zhang, B., Xiao, J., Jiao, J., Wei, Y., Zhao, Y.: Affinity attention graph neural network for weakly supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(11), 8082–8096 (2021)
- [106] Zhang, B., Xiao, J., Wei, Y., Sun, M., Huang, K.: Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12765–12772 (2020)
- [107] Zhang, B., Xiao, J., Wei, Y., Zhao, Y.: Credible dual-expert learning for weakly supervised semantic segmentation. *International Journal of Computer Vision* **131**(8), 1892–1908 (2023)
- [108] Zhang, D., Zhang, H., Tang, J., Hua, X.S., Sun, Q.: Causal intervention for weakly-supervised semantic segmentation. *Advances in Neural Information Processing Systems* **33**, 655–666 (2020)
- [109] Zhang, F., Gu, C., Zhang, C., Dai, Y.: Complementary patch for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7242–7251 (2021)
- [110] Zhang, X., Wang, Q., Zhang, J., Zhong, Z.: Adversarial autoaugment. arXiv preprint arXiv:1912.11188 (2019)
- [111] Zhou, T., Zhang, M., Zhao, F., Li, J.: Regional semantic contrast and aggregation for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4299–4309 (2022)
- [112] Zhou, Z.H.: A brief introduction to weakly supervised learning. *National science review* **5**(1), 44–53 (2018)
- [113] Zhu, Y., Zhou, Y., Xu, H., Ye, Q., Doermann, D., Jiao, J.: Learning instance activation maps for weakly supervised instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3116–3125 (2019)



Shaheer U. Saeed is currently a Research Fellow at the Department of Medical Physics and Biomedical Engineering, University College London. His current research interests include data-efficient learning within contexts of meta-learning, reinforcement learning, and deep learning, with applications to medical imaging.



Shiqi Huang is a PhD candidate from the School of Optics and Photonics, Beijing Institute of Technology, is currently undertaking a research visit at the Department of Medical Physics and Biomedical Engineering, University College London. Her primary research focus lies in medical image analysis, particularly in the areas of image segmentation and registration.



João Ramalinho is currently a Senior Research Associate at the Department of Medical Physics and Biomedical Engineering, University College London. His current research interests include techniques for intraoperative ultrasound guidance, and the application of machine learning to the field of surgical image guidance.



Iani J.M.B. Gayo is a final-year PhD candidate at the EPSRC Centre for Doctoral Training in Intelligent, Integrated Imaging In Healthcare (i4health), University College London. Her research focuses on applying reinforcement learning to assist with prostate cancer interventions.



Dean C. Barratt is currently a Professor at the Department of Medical Physics and Biomedical Engineering, University College London. His research is focused on developing software algorithms to analyse medical images and translating these into healthcare environments.



Nina Montaña-Brown is currently a PhD candidate at the Department of Medical Physics and Biomedical Engineering, University College London. She is also a Research Scientist at Verdure Imaging. Her current research interests include applications of deep learning to medical image analysis for image guided surgery.



Ester Bonmati is currently a Lecturer in Computer Science and Engineering at the University of Westminster and an Honorary Senior Research Fellow at the Department of Medical Physics and Biomedical Engineering, University College London. She specialises in the intersection of medical ultrasound and deep learning, making significant contributions to the diagnosis and treatment of cancer.



Matthew J. Clarkson is currently a Professor at the Department of Medical Physics and Biomedical Engineering, University College London. His research interests are focused on the development of systems for image-guided surgery, including image registration, computer graphics, visualisation, image processing and software architectures.



Stephen P. Pereira is currently a Professor of Hepatology and Gastroenterology at University College London. He is also a consultant at University College Hospital and the Royal Free Hospital. His clinical and research interests centre on the causes, early diagnosis and treatment of liver and pancreatic cancer, as well as benign biliary and pancreatic diseases.



Yipeng Hu is currently an Associate Professor at the Department of Medical Physics and Biomedical Engineering, University College London. His current research interest is applying machine learning and computational algorithms in medical image analysis and image-guided intervention.



Brian Davidson is currently Consultant HPB and Liver Transplant Surgeon at the Royal Free Hospital and Professor of Surgery at University College London. He is also the chair of the London NIHR RFPB grant panel. His grant income over the last 10 years has exceeded £31M. His research interests include evidence based healthcare (UK Cochrane Editor), clinical trials, pathogenesis of HPB cancers and novel surgical technologies.