

This document is supplementary material to “Outplaying Elite Table Tennis Players with an Autonomous Robot.”

Approximate Python Code

Here we provide approximate Python code that replicates the most critical aspects of our work. For the purposes of illustration, the code presented here is implemented in less computationally efficient ways than our actual implementation; has abstracted away many of the system implementation details such as remote process communication and check-pointing; and is written more directly than our actual implementation which supported wide configuration capabilities that otherwise detract from readability.

```

1 def train_loop():
2     models = Models(
3         policy=make_policy_network(),
4         af1=make_q_network(),
5         af2=make_q_network(),
6         af1_target=make_q_network(),
7         af2_target=make_q_network(),
8     )
9     # Sync prediction and target networks before start
10    copy_weights(models.af1.trainable_params, models.af1_target.trainable_params)
11    copy_weights(models.af2.trainable_params, models.af2_target.trainable_params)
12
13    opts = Optimizers(
14        policy=keras.optimizers.Adam(learning_rate=1.0e-4),
15        critic=keras.optimizers.Adam(learning_rate=1.0e-4),
16    )
17
18    param_server = start_parameter_server(models.policy)
19    replay_server = start_experience_replay_server()
20    table_datasets = [
21        make_dataset_for_table(replay_server, table_name)
22        for table_name in TABLE_NAMES
23    ]
24    task_server = start_task_factory_server()
25
26    # Wait for MINIMUM_STEPS to be recorded to replay before we start any training
27    replay_server.block_and_wait(MINIMUM_STEPS)
28
29    epoch = 0
30    while True:
31        for _ in range(BATCHES_PER_EPOCH):
32            transition_data = sample_data(table_datasets)
33            # see SAC code listing for implementation
34            losses = step_sac(models, opts, transition_data)
35            write_metrics(losses)
36            param_server.update_params(models.policy)
37            epoch += 1
38            # do a periodic evaluation of the policy
39            if epoch % EPOCHS_PER_EVAL == 0:
40                task_server.needs_eval = True
41
42 def sample_data(table_datasets: List[tf.data.Dataset]) -> Transition:
43     # For simplicity, this code omits handling the edge case when only a subset of
44     # the tables have data to sample.
45     # Before concat, each dataset has a different batch size based on the table
46     # weights
47     batches = [next(iter(dataset)) for dataset in table_datasets]
48     batches = nested_concat(batches, axis=0)
49
50     return Transition(
51         obs=batches.obs,
52         action=batches.action,
53         reward=batches.reward,
54         next_obs=batches.next_obs,
55         done=batches.done,
56     )

```

Listing 1: Training Loop Code

```

1 def step_sac(m: Models , opts: Optimizers , t: Transition) -> Losses:
2     critic_params = m.af1.trainable_params + m.af2.trainable_params
3     target_critic_params = m.af1_target.trainable_params + m.af2_target.
         trainable_params
4
5     # these values do not need gradients traced through them
6     policy_head_next = m.policy.policy_head(t.next_obs)
7     actions_next = m.policy.sample_from_head(policy_head_next)
8     log_prob_next = m.policy.log_prob(policy_head_next, actions_next)
9
10    # update policy
11    with tf.GradientTape() as tape:
12        policy_head = m.policy.policy_head(t.obs)
13        sampled_actions = m.policy.sample_from_head(policy_head)
14        log_prob = m.policy.log_prob(policy_head, sampled_actions)
15        policy_loss = policy_loss(
16            log_prob_samples=log_prob,
17            q1_sampled=m.af1.action_value(t.obs, sampled_actions),
18            q2_sampled=m.af2.action_value(t.obs, sampled_actions),
19        )
20    opts.policy.minimize(policy_loss, m.policy.trainable_params, tape)
21
22    # update critic networks
23    with tf.GradientTape() as tape:
24
25        critic_loss = critic_loss(
26            q1_observed=m.af1.action_value(t.obs, t.action),
27            q2_observed=m.af2.action_value(t.obs, t.action),
28            q1_sampled_next=m.af1_target.action_value(t.next_obs, actions_next),
29            q2_sampled_next=m.af2_target.action_value(t.next_obs, actions_next),
30            log_prob_next=log_prob_next,
31            reward=t.reward,
32            done=t.done,
33        )
34
35    opts.critic.minimize(critic_loss, critic_params, tape)
36
37    # move target network params toward prediction networks
38    for pred_param, target_param in zip(critic_params, target_critic_params):
39        target_param.assign(
40            SMOOTH_FACTOR * pred_param + (1.0 - SMOOTH_FACTOR) * target_param
41        )
42
43    return Losses(policy_loss, critic_loss)
44
45 def policy_loss(
46     log_prob_samples: tf.Tensor, # (N, 1)
47     q1_sampled: tf.Tensor, # (N, 1)
48     q2_sampled: tf.Tensor, # (N, 1)
49 ) -> tf.Tensor:
50     min_q = tf.minimum(q1_sampled, q2_sampled)
51     return tf.math.reduce_mean(ALPHA * log_prob_samples - min_q)
52
53 def td_target(reward: tf.Tensor, done: tf.Tensor, discount: tf.Tensor, next_v: tf.
         Tensor) -> tf.Tensor:
54     return reward + (1 - done) * discount * next_v
55
56 def critic_loss(
57     q1_observed: tf.Tensor, # (N, 1)
58     q2_observed: tf.Tensor, # (N, 1)

```

```

59     q1_sampled_next: tf.Tensor, # (N, 1)
60     q2_sampled_next: tf.Tensor, # (N, 1)
61     log_prob_next: tf.Tensor, # (N, 1)
62     reward: tf.Tensor, # (N, 1)
63     done: tf.Tensor, # (N, 1)
64 ):
65     min_q = tf.minimum(q1_sampled_next, q2_sampled_next)
66     q_target = tf.stop_gradient(td_target(reward, done, DISCOUNT, min_q - ALPHA *
67                                     log_prob_next))
68     q1_loss=tf.keras.losses.MeanSquaredError(q_target, q1_observed),
69     q2_loss=tf.keras.losses.MeanSquaredError(q_target, q2_observed),
70     return q1_loss + q2_loss

```

Listing 2: SAC Code

```

1 def get_task(context: DistributedTaskContext) -> Task:
2     # The returned task can be either an experience collection task, or an evaluation
3     task
4     task_server = get_task_server()
5     # The model_state contains information like env_steps, epoch_number, policy model
6     # such that the desired logic can be implemented in the tasks
7     model_state = context.model_state
8     if task_server.needs_eval:
9         # Time for periodic evaluation
10        task_server.needs_eval = False
11        # In evaluation task only the model_state is needed
12        return DistributedEvalTask(
13            model_state=model_state,
14        )
15    else:
16        # Standard collect task
17        # replay_configuration contains the configuration
18        # to send data to replay buffer
19        replay_configuration = context.replay_config
20        # In collect task both model_state and replay_configuration are needed
21        return DistributedCollectTask(
22            model_state=model_state,
23            replay_config=replay_configuration,
24        )

```

Listing 3: The task factory is run in a separate process and is remotely queried by rollout workers.

```

1 def rollout_loop(context: DistributedTaskContext):
2     faoc = make_faoc_controller()
3     env_server = get_environment_server()
4     replay_server = get_experience_replay_server()
5     while True:
6         task = get_task(context)
7         # Learning starts after collecting a min amount of transitions
8         # before that, use Ornstein-Uhlenbeck noise policy
9         if task.model_state.epoch.value == 0:
10            agent_policy = OrnsteinUhlenbeckPolicy()
11        else:
12            agent_policy = task.model_state.policy
13
14        env_server.launch_environment() # Start the environment
15        obs = env_server.get_first_obs()
16        p_obs = preprocess(obs) # Preprocess observation for policy input
17        done = False
18        while not done:
19            action = agent_policy(p_obs) # Get action from policy

```

```

20         reference_traj = faoc(action) # Get reference trajectory from FAOC
21         env_server.step(reference_traj) # Step env with robot ref traj
22         # Get the next observation 32ms later
23         next_obs = env_server.wait_for_next_action_intent()
24         reward = compute_reward(obs, next_obs) # step reward
25         done = env_server.check_termination_event() # Check if episode ended
26
27         if not task.is_eval:
28             transition = Transition(
29                 obs=p_obs,
30                 action=action,
31                 reward=reward,
32                 next_obs=preprocess(next_obs),
33                 done=done,
34             )
35             # Write transition to replay buffer
36             replay_server.write(transition, task.replay_config)
37
38             obs = next_obs
39             p_obs = preprocess(next_obs)
40
41         write_metrics(env_server.get_episode_metrics()) # Write episode metrics

```

Listing 4: Rollout worker code. 30 instances are run in parallel.

```

1 import pygad
2
3 def evaluation_wrapper(ga_instance: pygad.GA, solution: np.ndarray, solution_idx: int) -> float:
4     """Wrapper function for evaluating a genome"""
5     genes = genome_evaluation(solution)
6
7     # Get the serve trajectory from the genes and the MPC
8     solution = genes_to_solve(genes)
9     if solution is None:
10         return 0.0
11
12     # Run the physics on the robot trajectory and record the behavior of the ball
13     ball, robot, racket_contacts, robot_contacts, bounces, net_hits = simulate_solve(solution)
14
15     # If the serve is not legal, award a small fitness based on the illegality
16     legal, fitness = is_solve_legal(ball, robot, racket_contacts, robot_contacts, bounces,
17                                     net_hits)
18
19     if legal:
20         # Compute the fitness of the serve based on the overall trajectory and
21         # the ball state (pos, vel, spin) at the second bounce
22         fitness += compute_fitness(ball, ball[bounces[1]])
23
24     return fitness
25
26 def serve_training() -> np.ndarray:
27     gen_alg = pygad.GA(
28         # Tunable parameters
29         num_generations=100,
30         num_parents_mating=6,
31         mutation_num_genes=3,
32         stop_criteria="saturate_40",
33         # Fixed parameters
34         random_mutation_min_val=-0.20,
35         random_mutation_max_val=0.20,
36         num_genes=9,
37         fitness_func=evaluation_wrapper,
38         init_range_low=0.0,
39         init_range_high=1.0,
40         gene_space={"low": 0.0, "high": 1.0},
41     )
42     gen_alg.run()

```

```

42     genes, _, _ = gen_alg.best_solution()
43     return genes

```

Listing 5: Serve training pipeline

```

1  class BallContact(Enum):
2      AIR = 0 # ball is airborne, i.e., no contact at all
3      RACKET = 1 # ball contact with racket
4      TABLE = 2 # ball contact with table
5
6
7  class BallPosition:
8      """Timestamped external ball position obtained from the overhead perception system
9      triangulations"""
10
11     timestamp: float
12     x: float # x position [m]
13     y: float # y position [m]
14     z: float # z position [m]
15
16  class BallVelocity:
17      """Timestamped ball velocity estimated from successive ball positions in the ball state
18      estimator"""
19
20     timestamp: float
21     x: float # x velocity [m/s]
22     y: float # y velocity [m/s]
23     z: float # z velocity [m/s]
24
25  class GazeControlSystem:
26      """System consisting of pan/tilt galvanometer mirrors and an electrically tunable focal power
27      lens"""
28
29     pan_angle: float # pan angle of the vertical mirror
30     tilt_angle: float # tilt angle of the horizontal mirror
31     focal_power: float # focal power of the liquid lens
32
33     T_gcs_world: np.ndarray # homogeneous transform from world to gcs
34
35     def track_ball_position(self, position_world: BallPosition):
36         pass # set pan/tilt angles to track the ball and adjust focal power based on distance (
37         pre-calibrated)
38
39  class Event:
40      """Single definition for clarity; An event can either be in
41      * 2D pixel coordinates,
42      * 3D metric coordinates, or
43      * unit spherical coordinates"""
44
45     timestamp: float # in seconds
46     x: int | float # x-coordinate [pixel/unit]
47     y: int | float # y-coordinate [pixel/unit]
48     polarity: bool # on/off events
49
50     # optional
51     z: int | float # z-coordinate [pixel/unit]
52     polar: float # polar coordinate [rad]
53     azimuth: float # azimuth coordinate [rad]
54
55  class EventCamera:
56      # note: in practice, we use a 320x320 pixel center region of interest (ROI) for efficiency
57     width: int # sensor width [pixels]
58     height: int # sensor height [pixels]
59
60     polarities = [0] # we only accumulate the OFF polarity in practice
61
62     def generate_events(self) -> Generator[Event]:

```

```

63         while True:
64             yield Event()
65
66     def events_to_time_surface(self, events: list[Event]) -> np.ndarray:
67         """a time surface is a (per-polarity) image of the most recent event timestamp at every (
68         x,y) location"""
69         time_surface = np.zeros(self.height, self.width, len(self.polarities))
70         for timestamp, x, y, polarity in events:
71             time_surface[y, x, polarity] = timestamp
72         return time_surface
73
74     @classmethod
75     def time_surface_to_events(time_surface: np.ndarray) -> Generator[Event]:
76         width, height, polarities = time_surface.shape
77         for x in range(width):
78             for y in range(height):
79                 for polarity in polarities:
80                     timestamp = time_surface[y, x]
81                     if timestamp != 0:
82                         yield Event(timestamp=timestamp, x=x, y=y, polarity=polarity)
83
84 class BallDetection:
85     center_x: float # x-coordinate of the bounding circle center [pix]
86     center_y: float # x-coordinate of the bounding circle center [pix]
87     radius: float # radius of the bounding circle [pix]
88     confidence: float # confidence in the bounding circlce [0-1]
89
90
91 class BallDetectorCNN:
92
93     tensorrt_model: Callable[[list[np.ndarray]], list[BallDetection]] # trained with torch and
94     converted using trtexec
95
96     def detect_balls(self, time_surfaces: list[np.ndarray]) -> list[BallDetection]:
97         return self.tensorrt_model(time_surfaces)
98
99     def mask_ball(time_surface: np.ndarray, ball_detection: BallDetection):
100         """mask the time surface such that only events on the ball surface remain"""
101         x, y, radius, _ = ball_detection
102         mask = np.zeros_like(time_surface, dtype=bool)
103         circle_mask = cv.circle(mask, (x, y), radius, fill=True)
104         return time_surface & circle_mask
105
106 class SpinEstimate:
107     """Timestamped angular velocity (spin) estimate with associated uncertainty (covariance or
108     variance)"""
109
110     timestamp: float
111     spin: np.ndarray # 3d angular velocity
112
113     # depending on the spin esimation method (CNN or CMax) we have different measures of
114     uncertainty
115     covariance: float # heterescedastic uncertainty learned by the spin estimation CNN (lower =
116     better)
117     variance: float # variance/contrast of image of warped events spin spin refinement with CMax
118     (higher = better)
119
120     @property
121     def axis(self):
122         return self.spin / self.magnitude()
123
124     @property
125     def magnitude(self):
126         return np.linalg.norm(self.spin)
127
128 class SpinEstimatorCNN:
129     """CNN trained on pseudo ground-truth labels obtained with CMax that also predicts
130     heteroscedastic uncertainty"""

```

```

128 tensorrt_model: Callable[[list[np.ndarray]], list[SpinEstimate]] # trained with torch and
    converted using trtexec
129
130 def estimate_spins(self, time_surface: list[np.ndarray]) -> list[SpinEstimate]:
131     """batch inference"""
132     return self.tensorrt_model(time_surface)
133
134
135 class BallStateEstimator:
136     """Pseudo ball state estimator"""
137
138     position: BallPosition
139     velocity: np.ndarray
140     spin: SpinEstimate
141     ball_physics_model: Callable[[float, float], Any] # models air drag, gravity, etc.
142
143     # the main loop needs to know if a ball contact occurred; we model this with callbacks
144     detect_ball_contact: Callable[[BallPosition, BallVelocity], BallContact] # detect if ball
    contact occurred
145     contact_callback: Callable[[BallContact]] # notify caller about the ball contact
146
147     # the latency is estimated using the difference in hardware trigger and triangulation arrival
    timestamp
148     current_latency: float
149
150     def predict(self, timestamp: float) -> Any:
151         """predict the state of the ball given a time point"""
152         ball_contact = self.detect_ball_contact(self.position, self.velocity)
153         self.contact_callback(ball_contact)
154         if ball_contact == BallContact.TABLE:
155             ... # apply table contact model to state variables to predict post-contact spin
    estimate
156         elif ball_contact == BallContact.RACKET:
157             self.reset()
158         elif ball_contact == BallContact.AIR:
159             self.ball_physics_model(timestamp, self.current_latency)
160             ...
161         return self.position, self.spin
162
163     def update(self, observation: BallPosition | SpinEstimate) -> Any:
164         """update the filter with a new observation (and its associated uncertainty)"""
165         self.current_latency = observation.timestamp - ...
166         ...
167         return self.position, self.spin
168
169     def reset(self):
170         """reset filter"""
171         ...
172
173     def register_contact_callback(self, callback: Callable[[BallContact]]):
174         self.contact_callback = callback
175
176
177 class MedianFilter:
178     """Simple median filter to reject false ball detections"""
179
180     window_size: int
181     observations: list[Any]
182
183     def filter(self, observation: Any) -> Any:
184         """median filter the observation and immediately return the result"""
185         # note: if observation is list-like, the median is computed per component
186         if len(self.observations) > self.window_size:
187             self.observations.pop(0)
188             self.observations.append(observation)
189         return np.median(self.observations)
190
191     def reset(self):
192         self.observations = []
193
194
195 class ContrastMaximization:

```



```

196 """Note: iterative version for clarity; we have implemented this on the GPU"""
197
198 def events2d_to_3d(events2d: list[Event], ball_detection: BallDetection) -> Generator[Event]:
199     """Convert 2d events to 3d events on the unit sphere using the ball detection (
    orthographic projection)"""
200     center_x, center_y, radius, _ = ball_detection
201     for event2d in events2d:
202         timestamp, x2d, y2d, polarity = event2d
203         x3d, y3d = ((x2d, y2d) - (center_x, center_y)) / radius
204         z3d = -np.sqrt(1 - (x3d**2 + y3d**2))
205         yield Event(timestamp=timestamp, x=x3d, y=y3d, z=z3d, polarity=polarity)
206
207 def warp_events3d_by_spin(events3d: list[Event], spin_estimate: SpinEstimate) -> Generator[
    Event]:
208     """Rotate 3d events to a reference timestamp using a candidate spin estimate"""
209     timestamp_ref = events3d[-1].timestamp
210     for event3d in events3d:
211         timestamp, x, y, z, polarity = event3d
212         td = event3d.timestamp - timestamp_ref
213         rotation = Rotation.from_rotvec(-spin_estimate * td)
214         x_rot, y_rot, z_rot = rotation.apply((x, y, z))
215         yield Event(timestamp=timestamp, x=x_rot, y=y_rot, z=z_rot, polariy=polarity)
216
217 def event3d_to_spherical(events3d: list[Event]) -> Generator[Event]:
218     for event3d in events3d:
219         timestamp, x, y, z, polarity = event3d
220         polar, azimuth = np.arcsin(z), np.arctan2(y, x)
221         yield Event(timestamp=timestamp, polar=polar, azimuth=azimuth, polarity=polarity)
222
223 def spherical_to_image_of_warped_events(events_spherical: list[Event], size: tuple[int, int])
    -> np.ndarray:
224     """Generate an image of warped events in spherical coordinates"""
225     height, width = size
226     image_of_warped_events = np.zeros((height, width))
227     for event_spherical in events_spherical:
228         polar, azimuth = event_spherical
229         x = width * (azimuth / (2 * np.pi) + 0.5)
230         y = height * ((0.5 - polar) / np.pi)
231         image_of_warped_events[y, x] += 1
232     return image_of_warped_events
233
234
235 class SpinRefinerCMax:
236     """Returns the spin candidate that best explains the ball rotation by contrast-maximizing (
    CMax) the image of warped events"""
237
238     contrast_maximizer: ContrastMaximization
239
240     def refine_spin_estimate(
241         self, events2d: list[Event], spin_candidates: SpinEstimate, ball_detection: BallDetection
242     ) -> SpinEstimate:
243
244         # note: in practice, this for loop is executed in parallel on the GPU (for all spin
    candidates simultaneously)
245         best_spin_candidate = SpinEstimate()
246         max_variance = 0
247         for spin_candidate in spin_candidates:
248             events3d = self.contrast_maximizer.events2d_to_3d(events2d, ball_detection)
249             events3d_warped = self.contrast_maximizer.warp_events3d_by_spin(events3d,
    spin_candidate)
250             events_spherical = self.contrast_maximizer.event3d_to_spherical(events3d_warped)
251             image_of_warped_events = self.contrast_maximizer.spherical_to_image_of_warped_events(
    events_spherical)
252             variance = np.var(image_of_warped_events)
253             if variance > max_variance:
254                 best_spin_candidate = spin_candidate
255                 best_spin_candidate.variance = variance
256
257         return best_spin_candidate
258
259
260 class BallTrackingAndSpinEstimation:

```

```

261 """Main class modeling the Gaze Control System operation: ball tracking and spin estimation
262 """
263 gaze_control_systems: list[GazeControlSystem]
264 event_cameras: list[EventCamera]
265 ball_state_estimator: BallStateEstimator
266
267 # time surfaces
268 time_surfaces: list[np.ndarray] # shared between threads
269
270 # CNN-based ball detection
271 ball_detector_cnn: BallDetectorCNN
272 ball_detections_filtered: list[BallDetection] # shared from ball detection -> spin
273 estimation thread
274
275 # CNN-based spin estimation
276 spin_estimator_cnn: SpinEstimatorCNN
277 best_spin_estimate: SpinEstimate # shared from spin estimation -> spin refinement thread
278 best_spin_index: int # shared from spin estimation -> spin refinement thread
279
280 # CMax-based spin refinement
281 spin_refiner_cmax: SpinRefinerCMax
282 refined_spin_estimate_cmax: SpinEstimate # shared from spin refinement -> spin estimation
283 thread
284
285 # publishers/subscribers
286 triangulation_subscriber: Subscriber
287 spin_estimation_publisher: Publisher
288
289 def init(self, num_systems=3):
290     self.gaze_control_systems = [GazeControlSystem() for _ in range(num_systems)]
291     self.median_filters = [MedianFilter(window_size=5) for _ in range(num_systems)]
292
293     # start all threads
294     self.should_exit = False # set by user interaction, e.g. Ctrl+C
295
296     # initialize spin estimates
297     self.best_spin_estimate = None
298     self.refined_spin_estimate_cmax = None
299
300     # register a contact callback (e.g., ball with racket or table) to reset spin estimates
301     self.ball_state_estimator.register_contact_callback(self.contact_callback)
302
303     # the ball triangulations are obtained from the overhead perception system at 200Hz in a
304     # callback
305     # the ball position is filtered and extrapolated in time which allows for smooth as-fast-
306     # as-possible tracking
307     self.triangulation_subscriber.callback(self.ball_triangulation_callback)
308     self.ball_tracking_thread()
309
310     # 1) detect ball, 2) estimate spin, and 3) refine spin (asynchronously but with
311     # dependence on each other)
312     self.ball_detection_thread()
313     self.spin_estimation_thread()
314     self.spin_refinement_thread()
315
316 @callback
317 def contact_callback(self, ball_contact: BallContact):
318     # reset spin estimates when racket or table contact is detected
319     if ball_contact is BallContact.RACKET or BallContact.TABLE:
320         self.best_spin_estimate = None
321         self.refined_spin_estimate_cmax = None
322
323 @callback
324 def ball_triangulation_callback(self, ball_position: BallPosition):
325     # estimate ball state (and latency); in practice, this is called at 5ms intervals (200Hz)
326     self.ball_state_estimator.update(position=ball_position)
327
328 @thread
329 def ball_tracking_thread(self):
330     """consumes: ball state estimate, produces: per-system mirror/lens control commands"""

```

```

327         while not self.should_exit:
328
329             # predict the current state of the ball (accounting for triangulation latency and
330             # aerodynamics)
331             ball_position_predicted, _ = self.ball_state_estimator.predict(time.now())
332
333             # track the ball with each gaze control system
334             (gcs.track_ball_position(ball_position_predicted) for gcs in self.
335             gaze_control_systems)
336
337     @thread
338     def ball_detection_thread(self):
339         """consumes: per-system time surfaces, produces: per-system filtered ball detections"""
340         while not self.should_exit:
341             # note: in practice, we accumulate events directly into the time surface for
342             # efficiency
343             self.time_surfaces = [
344                 event_camera.events_to_time_surface(event_camera.generate_events())
345                 for event_camera in self.event_cameras
346             ]
347
348             # detect ball on time surfaces using batched inference
349             ball_detections = self.ball_detector_cnn.detect_balls(self.time_surfaces)
350
351             # apply median filter to ball detections
352             self.ball_detections_filtered = [
353                 median_filter(ball_detection)
354                 for median_filter, ball_detection in zip(self.median_filters, ball_detections)
355             ]
356
357     @thread
358     def spin_estimation_thread(self):
359         """consumes: per-system time surfaces and ball detections, produces: single (possibly
360         refined) spin estimate"""
361         while not self.should_exit:
362
363             # mask events outside of the ball bounding circle
364             time_surfaces_masked = [
365                 BallDetectorCNN.mask_ball(time_surface, ball_detection)
366                 for time_surface, ball_detection in zip(self.time_surfaces, self.
367                 ball_detections_filtered)
368             ]
369
370             # estimate spin on masked time surfaces using batched inference
371             spin_estimates = self.spin_estimator_cnn.estimate_spins(time_surfaces_masked)
372
373             # best spin estimate has lowest covariance/uncertainty
374             self.best_spin_index = np.argmin(spin_estimates, key=lambda spin_estimate:
375             spin_estimate.covariance)
376             best_spin_estimate = spin_estimates[self.best_spin_index]
377
378             # filter best spin estimate before making it available for refinement
379             _, self.best_spin_estimate = self.ball_state_estimator.update(spin=best_spin_estimate
380             )
381
382             # if available with low uncertainty (= high variance), overwrite the CNN's best spin
383             # estimate magnitude
384             # we assume that the spin axis predicted by the CNN is already accurate, thus only
385             # refine the magnitude
386             MIN_VARIANCE = 0.01
387             if self.refined_spin_estimate_cmax is not None and self.refined_spin_estimate_cmax.
388             variance > MIN_VARIANCE:
389                 self.best_spin_estimate.magnitude = self.refined_spin_estimate_cmax.magnitude
390
391             # publish for robot control
392             self.spin_estimation_publisher.publish(self.best_spin_estimate)
393
394     @thread
395     def spin_refinement_thread(self):
396         """consumes: single spin estimate, produces: single refined spin estimate"""
397         while not self.should_exit:

```

```

389         # only refine once we have a spin estimate with high enough magnitude
390         MIN_SPIN_MAGNITUDE = 50 # rad/s
391         if self.best_spin_estimate is None or self.best_spin_estimate.magnitude <
MIN_SPIN_MAGNITUDE:
392             continue
393
394         # generate N magnitude-varying spin candidates around the current best estimate
395         MAX_ERROR_PERC, MAX_SPIN_MAGNITUDE, MAX_NUM_SPIN_CANDIDATES = 0.7, 1000, 1000
396         magnitude = self.best_spin_estimate.magnitude()
397         mag_lower_bound = max(0, magnitude * (1 - MAX_ERROR_PERC))
398         mag_upper_bound = min(magnitude * (1 + MAX_ERROR_PERC), MAX_SPIN_MAGNITUDE)
399         magnitudes = np.linspace(mag_lower_bound, mag_upper_bound, MAX_NUM_SPIN_CANDIDATES)
400         spin_candidates = magnitudes * self.best_spin_estimate.axis()
401
402         # convert back to events representation for refinement with contrast maximization
403         events: list[Event] = EventCamera.time_surface_to_events(self.time_surfaces[self.
best_spin_index])
404
405         # evaluate all spin candidates in parallel on the GPU for efficiency
406         # the refined spin estimate is the one that maximizes the variance of the image of
warped events
407         self.refined_spin_estimate_cmax = self.spin_refiner_cmax.refine_spin_estimate(events,
spin_candidates)
408
409
410 def main():
411     BallTrackingAndSpinEstimation().init(num_systems=3) # run the gaze control system(s)

```

Listing 6: Gaze Control System pseudocode