Robotic Technologies for Solar-Powered UAVs: Fully-Autonomous Updraft-Aware Aerial Sensing for Multi-Day Search-and-Rescue Missions

Philipp Oettershagen, Thomas Stastny, Timo Hinzmann, Konrad Rudin, Thomas Mantel, Amir Melzer, Bartosz Wawrzacz, Gregory Hitz and Roland Siegwart

Autonomous Systems Lab
Swiss Federal Institute of Technology Zurich (ETH Zurich)
Leonhardstrasse 21
8092 Zurich
+41 44 632 7395
philipp.oettershagen@mavt.ethz.ch

Abstract

Large-scale aerial sensing missions can greatly benefit from the *perpetual endurance* capability provided by high-performance low-altitude solar-powered UAVs. However, today these UAVs suffer from small payload capacity, low energetic margins and high operational complexity. To tackle these problems, this paper presents four individual technical contributions and integrates them into an existing solar-powered UAV system: First, a lightweight and power-efficient day/night-capable sensing system is discussed. Second, means to optimize the UAV platform to the specific payload and to thereby achieve sufficient energetic margins for day/night-flight with payload are presented. Third, existing autonomous launch and landing functionality is extended for solar-powered UAVs. Fourth, as a main contribution an extended Kalman filter-based autonomous thermal updraft tracking framework is developed. Its novelty is that it allows the end-to-end integration of the thermal-induced roll moment into the estimation process. It is assessed against unscented Kalman filter and particle filter methods in simulation and implemented on the aircraft's low-power autopilot. The complete system is verified during a 26-hour search-and-rescue aerial sensing mockup mission that represents the first-ever fully-autonomous *perpetual endurance* flight of a small solar-powered UAV with a day/night-capable sensing payload. It also represents the first time that solar-electric propulsion and autonomous thermal updraft tracking are combined in flight. In contrast to previous work that has focused on the energetic feasibility of *perpetual flight*, the individual technical contributions of this paper are considered core functionality to guarantee ease-of-use, effectivity and reliability in future multi-day aerial sensing operations with small solar-powered UAVs.

1 Introduction

1.1 Solar-Powered Unmanned Aerial Vehicles: State of the Art and Challenges

Solar-powered fixed-wing Unmanned Aerial Vehicles (UAVs) promise significantly increased flight endurance over aerial vehicles powered by electric batteries or gas alone. Large-scale disaster relief, meteorological surveys in remote areas and continuous border or maritime patrol applications benefit in particular from the multi-hour or even multi-day continuous flight capability provided by these robotic systems (Colella & Wenneker, 1996). In contrast to their large-scale High-Altitude Long-Endurance (HALE) UAV counterpart (Ackerman, 2013), small-scale solar-powered UAVs are mostly designed for Low-Altitude Long-Endurance (LALE) applications. They have to deal with the more challenging meteorology of the lower atmosphere (clouds, rain, wind gusts and thermals), but provide the advantage of higher resolution imaging with reduced cloud obstruction at lower overall system complexity. Motivated by recent advances in battery- and solar-technology, solar-powered non-perpetual-flight capable LALE UAVs with targeted flight times of up to 14 hours have been studied (Weider et al., 2007; Malaver, Gonzalez, Motta, & Villa, 2015) and are currently being intensely investigated by the industry (AeroVironment, 2013; ByeAerospace, 2015).

Figure 1: The aircraft- and robotic-technologies developed to achieve the first-ever fully-autonomous multi-day solar-powered Search-and-Rescue (SAR) flight: Autonomous launch of the *AtlantikSolar AS-3* UAV (top right), use of automated thermal updraft tracking and solar-powered flight to replenish the batteries while performing a large-scale victim search mission using the RGB+Infrared camera payload (left), continuation of the search-and-rescue mission under battery power during the night (middle- and bottom-right), and autonomous landing (not shown) after a 26-hour flight without a single pilot intervention.

However, the deployment of low-altitude solar-powered UAVs which are energy-wise capable of *perpetual flight* to real-world missions has been hindered by the operational complexities involved: First, the requirement for low power consumption but significant battery and solar-module carrying capacity results in slow, undamped and coupled flight dynamics that require very skilled pilots. Second, while the current performance of solar- and battery-technology has resulted in demonstrations of *perpetual flight*, it has not allowed *perpetual flight* with a meaningful sensing payload at low altitude yet. More specifically, the first-ever solar-powered multi-day flight was a 48-hour flight with the 12 kg *SoLong* platform by Cocconi (2005): Eight pilots were required to perform the flight, thermal udprafts needed to be tracked manually to restrain the power expenditure, no payload was carried except for an RGB video-feed used for control by the pilots, and the aircraft crossed the night with only two hours of reserve power in the batteries. A 27-hour flight with the 2.5 kg *SkySailor* UAV was demonstrated by Noth (2008): The aircraft flew autonomously except for launch and landing, carried no payload, and only had 5.8 % of battery energy remaining at the end of the night. More recently, we have presented the 81-hour continuous flight of the 7 kg *AtlantikSolar AS-2* UAV — the current unofficial flight endurance *world record* for all aircraft below 50 kg total mass — in our own work (Oettershagen et al.,

2017). The UAV operated autonomously except for launch and landing and had 39 % of reserve power (6.8 h of reserve time) remaining. However, it also relied on skilled pilots for the challenging launch and landing and did not carry any payload although the energetic margins in hindsight indicate that this was feasible. Clearly, today's challenge in low-altitude solar-powered perpetual flight lies in transferring the technology from the research stage into real-life missions. In that context, UAV operators and end-users such as search-and-rescue teams will require systems that, first, guarantee reliable day/night-operation with a meaningful sensing payload and, second, provide ease-of-use through full autonomy from launch to landing. The robotic technologies presented in this paper are developed to answer this end-user need.

1.2 Contributions

Our previous work (Oettershagen et al., 2017) has focused on the platform design and low-level flight autonomy required for energetically-reliable solar-powered perpetual flight. The goal of the research presented in this paper is to increase the ease-of-use and effectiveness of such solar-powered UAVs in *real-life* aerial sensing missions. Therefore, we extend the existing low-level autonomy with more advanced robotic capabilities that are directly motivated by the goal to achieve the first-ever fully-autonomous energetically-perpetual flight with a day/night-capable sensing payload on a small solar-powered UAV. This goal was achieved with a 26-hour search-and-rescue mockup mission — which did not require a single pilot intervention — performed from July 19–20th 2016 by the *AtlantikSolar AS-3* UAV (Figure 1). Results and field deployment experience from this flight are presented.[1] The individual research- and technological contributions that were motivated by and developed for this specific search-and-rescue flight, but with the greater goal of increasing both ease-of-use and effectiveness of solar-powered UAVs in similar applications, are:

- *Day/Night-capable SAR-payload with automatic human-detection capability*: The development of a lightweight, power-efficient yet day/night-capable search-and-rescue payload consisting of a color camera, an infrared (IR) camera, an onboard computer and WLAN for persistent data downlink is presented. Results from applying our existing human detection algorithms (Kümmerle, Hinzmann, Siegwart, & Gilitschenski, 2016) during the 26-hour SAR flight are discussed.
- *Platform optimization for perpetual flight with the SAR payload*: Our previous work (Oettershagen et al., 2017) suggested the feasibility of perpetual flight with the aforementioned payload and assessed the theory behind adapting the solar-powered UAV's energy-storage system to the specific payload weight and power consumption. This paper applies and verifies these methods using the 26-hour flight.
- *Autonomous launch and landing framework*: The existing control structure is extended with an automatic launch and landing framework that addresses the fragility of solar-powered UAV platforms through preventive features (e.g. flare before touchdown) and a hierarchical chain of safety checks and contingency measures (e.g. go-around on wrong landing altitude).
- *Autonomous thermal updraft tracking framework*: As the paper's main theoretical contribution, a framework to optimize the power consumption by automated tracking of thermal updrafts is developed. First, the framework extends previous approaches through a better sink rate model. Second, it improves the estimation problem's observability by introducing a second measurement (the roll moment induced by a thermal) into the Kalman filtering process. Our 26-hour mission also yields the first combined flight results of solar-electric propulsion and autonomous thermal updraft tracking in the literature.

The remainder of this paper is organized as follows. Section 2 presents the design of the baseline *AtlantikSolar* UAV platform. Section 3 introduces the theoretical background and presents first results for the platform, sensing system and flight autonomy extensions implemented in this paper. Section 4 presents field deployment results from the fully autonomous 26-hour search-and-rescue flight and analyzes the efficacy of each individual technological contribution. Concluding remarks are provided in Section 5.

[1] A video of the flight is available at https://www.youtube.com/watch?v=8m76Mx9m2nM.

2 System Baseline: The Solar-Powered AtlantikSolar UAV

The *AtlantikSolar* UAV is a solar-powered robotic aircraft designed for multi-day continuous flight. Initial design and flight results were presented in (Oettershagen, Melzer, Mantel, Rudin, Lotz, et al., 2015), and first aerial sensing missions with the UAV were described in (Oettershagen, Stastny, et al., 2015). The most comprehensive description of the UAV's conceptual design, detailed design, and flight results — including an 81-hour continuous solar-powered flight that represents the current *world record* in flight endurance for all aircraft below 50 kg total mass — is provided in (Oettershagen et al., 2017). The reader is referred to the above literature for an in-depth design review. The following section briefly summarizes this *baseline* UAV design, i.e. the airframe and control system of the *AtlantikSolar AS-2* UAV used for the 81-hour flight. Section 3 then describes the specific optimizations implemented in this paper to achieve a fully-autonomous 26-hour search-and-rescue flight using *AtlantikSolar AS-3*.

UAV Platform Hardware

The *AtlantikSolar* UAV airframe (Figures 1 and 2) is of a conventional glider-like T-tail configuration with two ailerons, an all-moving elevator and a rudder. The wing has a span of $B = 5.69\,\mathrm{m}$, a chord of $c_{\mathrm{wing}} = 0.305\,\mathrm{m}$, and consists of three pieces of similar span that can be disassembled before transport. It is perfectly rectangular, i.e. neither swept nor tapered, to house the two rows of solar cells over the whole wing span. A dihedral angle of $6\,^{\circ}$ on the outer wings provides basic roll axis eigenstability. The UAV's energy generation and storage system incorporates 88 SunPower E60 solar cells that provide $\eta_{\mathrm{sm}} = 23.7\%$ measured module-level efficiency. Energy storage is handled by cylindrical Lithium-Ion batteries (Panasonic NCR18650b with $251\,{}^{\mathrm{Wh}}/_{\mathrm{kg}}$ energetic density) that are fitted into the cylindrical wing spars to optimally distribute the mass in a *span loader* concept. For the *AtlantikSolar AS-2* UAV, the 60 Lithium-Ion cells provide $E_{\mathrm{bat}}^{\mathrm{max}} = 733\,\mathrm{Wh}$ at $m_{\mathrm{bat}} = 2.92\,\mathrm{kg}$ such that the total aircraft mass is $m_{\mathrm{tot}}^{\mathrm{AS\text{-}2}} = 6.93\,\mathrm{kg}$. The resulting stall speed is $v_{\mathrm{air}}^{\mathrm{stall}} = 7.4\,{}^{\mathrm{m}}/_{\mathrm{s}}$. The actuation system consists of four Volz DA 15-N servos that have successfully completed 30 days of continuous operation in reliability bench testing. The propulsion system contains a RS-E Strecker 260.20 brushless DC motor, a Kontronik Koby 55 LV motor controller, an all-steel planetary gearbox with a 5:1 reduction ratio and a foldable carbon-fiber propeller with diameter $D = 0.66\,\mathrm{m}$ and pitch $H = 0.6\,\mathrm{m}$. The avionics subsystem (Figure 2) is centered around a *Pixhawk* autopilot (Meier, Honegger, & Pollefeys, 2015) — an open-source software and open-source hardware project initiated at ETH Zurich — running a real-time operating system and featuring a 168 MHz Cortex M4F microprocessor with 256 KB RAM. For attitude estimation, an ADIS16448 Inertial Measurement Unit (IMU), a u-blox LEA-6H GPS receiver, and a Sensirion SDP600 differential pressure sensor are used. A 433 MHz medium-range telemetry link is integrated. The airplane implements a fully manual RC-command fall-back mode to deal with a severe case of autopilot failure. Night operations are possible due to four on-board indicator LEDs.

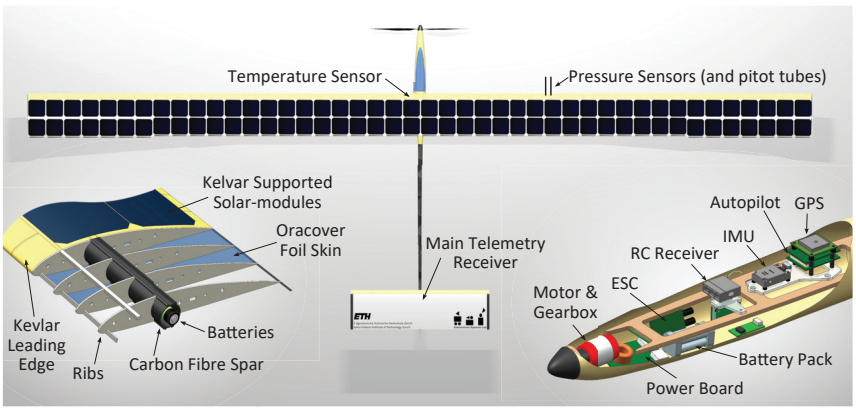

Figure 2: *AtlantikSolar* UAV airframe and avionics.

UAV Flight Control Architecture

The *AtlantikSolar* UAV flight control architecture emphasizes simplicity, robustness and low-power consumption to fulfill the need for reliable long-endurance robotic flight. The first stage of our flight control system employs a lightweight extended Kalman filter (EKF) to generate a drift-free aircraft state estimate. The EKF fuses data from the IMU with the magnetometer, GPS-position, GPS-velocity and airspeed measurements to successively estimate position, velocity, orientation, mean sea level static pressure (QFF) as well as accelerometer and gyroscope biases. The system achieves robustness against temporal GPS failure through the inclusion of airspeed measurements. Further details about the attitude estimation approach are provided in (Leutenegger, Melzer, Alexis, & Siegwart, 2014).

The custom-designed flight controller (Figure 3) features autonomous navigation including loitering and tracking of user-defined waypoints. For inner-loop control, our baseline-solution is a set of cascaded and saturated PID controllers: The Stability Augmentation System (SAS) applies rate-damping to shape the airplane's frequency response, while the Control Augmentation System (CAS) applies proportional-integral feedback to achieve roll (ϕ) and pitch (θ) reference tracking. Due to the specific design characteristics of solar-powered UAVs (high aspect ratio and thus high inertia especially in I_{xx} and I_{zz}, limited admissible structural loads) the flight controller implements numerous extensions to a common PID-based architecture: First, precise aircraft trim curves for the elevator ($u_{\text{ele}}^{\text{trim}} = f(v_{\text{air}})$), rudder ($u_{\text{rud}}^{\text{trim}} = f(v_{\text{air}})$) and aileron ($u_{\text{ail}}^{\text{trim}} = f(v_{\text{air}}, \phi)$) reduce tracking offsets. Second, coordinated turn control smooths the adverse yaw behavior and tracks the no-sideslip yaw (in body coordinates) rate $r = \frac{g \cdot \sin(\phi)}{v_{\text{air}}}$. Third, an overspeed protection mechanism is implemented. Fourth, a dynamic overload protection constantly monitors the vertical aircraft acceleration n_z and limits the elevator command dynamically to guarantee that $n_z < n_z^{\text{lim}}$.

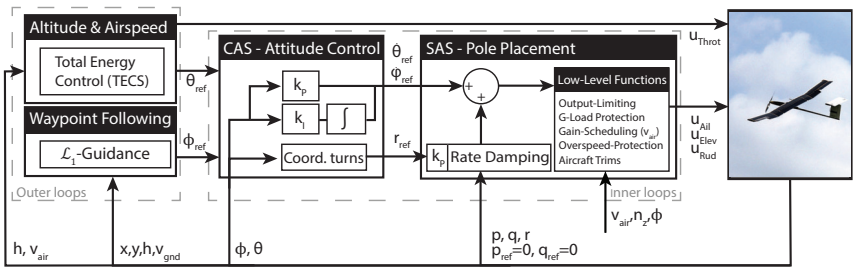

Figure 3: Cascaded control scheme implemented for ETH Zurich's solar-powered UAVs.

In its outer loop, the flight controller employs a nonlinear guidance law to track waypoints by generating the roll angle reference ϕ_{ref}. The process uses the current ground speed and heading along with a look-ahead distance $\mathcal{L}_1$ that is adapted online as outlined in (Park, Deyst, & How, 2007). Altitude and airspeed control are provided by an extended version of the *Pixhawk* (Meier et al., 2015) Total Energy Control System (TECS). The most important extension for the thermal updraft tracking implemented in this paper is the *thermal compliance*: In an updraft, the standard TECS implementation will decrease the pitch reference θ_{ref} to decrease the altitude if $h > h_{\text{ref}}$. Instead of actively working against thermals, we allow the UAV to gain potential energy from an updraft: Our TECS variant is configured such that θ_{ref} is fully and only used for airspeed control and u_{thr} only for altitude control. When at $h > h_{\text{ref}}$, the plane will thus choose $\theta_{\text{ref}}(t)$ such that $v_{\text{air}}(t) = v_{\text{air,ref}}(t)$ and will gradually reduce u_{thr}, potentially gaining altitude for strong thermals. When the aircraft has ascended such that $h \geq h_{\text{max}}$, the controller automatically engages the spoilers (upwards deflected ailerons) and gradually commands a pitch-down and thus maximum descent rate. This feature has proven to be of significant importance to avoid fly-aways and to thus increase airplane safety when operating in thermal updrafts (see Section 4). Note that our previous work (Oettershagen et al., 2017) contains a more in-depth description and verification of the flight control system.

3 Robotic Technologies for Fully-Autonomous Perpetual-Endurance Aerial Sensing

Fully-autonomous *perpetual-endurance* search-and-rescue support missions with a small solar-powered UAV require novel contributions in platform design, sensing system design and robotics. The central challenge lies in flying through the night with the additional power consumption of a day/night-capable SAR-payload. The development of an effective yet lightweight and power-efficient SAR-payload and energetic optimizations of the solar-powered UAV platform are thus presented in Section 3.1. The extended autonomous launch and landing functionality is presented in Section 3.2 and the autonomous thermal updraft tracking algorithm is discussed in Section 3.3. Together, these individual contributions result in novel solar-powered UAV technology that can be used more easily, reliably and effectively in large-scale aerial sensing applications.

3.1 UAV Platform and Sensing System

3.1.1 Payload Development

The specific purpose of the payload developed for this paper is to allow the real-time localization of victims from the altitude $h_{\text{nom}} \approx 100\,\text{m}$ above ground during both day and night of multi-day SAR flights. To this end, the payload integrates both a color camera and long-wavelength infrared (LWIR) camera with on-board image processing capabilities, but is at the same time designed to be lightweight and energy efficient. The camera types and optical characteristics chosen to fulfill these challenging constraints are described in Table 1. The focal lengths are selected to yield the same Field of View (FoV) for both cameras and the mounting angle allows both human detection applications and aerial mapping tasks. For optimal victim detection performance, the camera outputs are synchronously triggered and time-stamped by an FPGA-based visual-inertial sensing system (Nikolic et al., 2014). The two image streams are processed and stored on an onboard computer (UpBoard with an Intel Atom x5-Z8350 processor running ROS (Quigley et al., 2009)) and transmitted to the ground station using wireless LAN. The whole sensing system weighs $m_{\text{pld}} = 0.35\,\text{kg}$ and is fully integrated into the airframe to avoid additional aerodynamic drag. The measured payload power consumption during the day (with both cameras active) is 10.5 W. The payload integrates electronics to switch off the RGB camera during the night, thereby reducing the power consumption to 9.6 W. The payload can be monitored and, in case of a failure, hard-reset over *AtlantikSolar*'s telemetry link.

Table 1: Camera characteristics of the developed search-and-rescue sensing payload.

	Infrared camera	Color camera
Manufacturer	FLIR	MatrixVision
Type	Tau 2	mvBlueFOX3-M1020
Resolution	640x512 px	1600x1200 px
Focal length	19 mm	12 mm
Ground resolution	92 mm/px	39 mm/px
Field of view	32 ° x 25 ° (approximate)	
Mounting angle	75 ° (from horizontal plane)	
Triggering	1.5 Hz (synchronized)	

3.1.2 UAV Platform Optimization

The perpetual-flight performance of solar-powered UAVs is assessed and optimized in terms of the minimum battery state-of-charge SoC_{min}, the excess time T_{exc} (the reserve flying time guaranteed solely through battery energy at the end of the night) and the charge margin T_{cm} (the time with sunshine that is not required for charging because the batteries are already full). Clearly, these performance metrics and thus the overall

feasibility of perpetual flight depend on the payload. Carrying the $m_{\text{pld}} = 0.35\,\text{kg}$ of the aforementioned SAR payload and especially providing the $P_{\text{pld}} \approx 10\,\text{W}$ for the image processing during the whole night is a notable challenge for a small-scale solar-powered UAV that has to fly at constant altitude.

Figure 4 shows the payload-specific optimization performed in this paper for *AtlantikSolar AS-3* based on the conceptual design and analysis framework presented in (Oettershagen et al., 2017). Without payload, we expect a minimum state-of-charge of 39 % at June 21st and more generally a perpetual-flight capability until late September at $m_{\text{bat}} = 2.92\,\text{kg}$. With the SAR payload, T_{exc} and $SoC_{\text{min}} = 18\,\%$ are greatly decreased while T_{cm} decreases only slightly. With the existing platform configuration (*AtlantikSolar AS-2*), the energetic margins for perpetual-flight with this additional payload are critically low. Applying the design optimization however shows that at $m_{\text{bat}} = 3.8\,\text{kg}$ we have reached an equilibrium between T_{exc} and T_{cm} again and retrieve $SoC_{\text{min}} = 27.5\,\%$. Considering practical battery sizing constraints[2] we chose $m_{\text{bat}} = 3.5\,\text{kg}$ for *AtlantikSolar AS-3*. This configuration yields $SoC_{\text{min}} = 24.9\,\%$ on June 21st, and still yields $SoC_{\text{min}} = 21.6\,\%$, $T_{\text{exc}} = 3.0\,\text{h}$ and $T_{\text{cm}} = 3.87\,\text{h}$ on July 20th when the 26-hour flight was performed (Section 4.1). The additional payload mass and the 72 instead of 60 Li-Ion battery cells increase the overall system mass from $m_{\text{tot}}^{\text{AS-2}} = 6.92\,\text{kg}$ to $m_{\text{tot}}^{\text{AS-3}} = 7.73\,\text{kg}$ despite additional mass optimizations (lighter power electronics, redundancy battery pack and avionics) on *AS-3*. Consequently, the stall speed increases to $v_{\text{air}}^{\text{stall}} \approx 8.0\,\text{m/s}$.

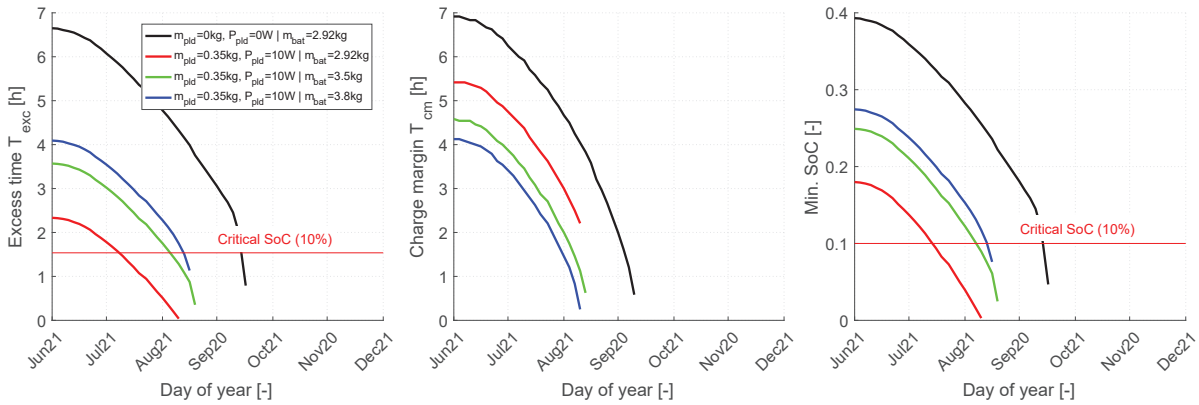

Figure 4: Expected excess time, charge margin and minimum battery state-of-charge for different payload- and solar-powered UAV platform configurations at different days of the year and a latitude of $\Phi = 47°N$. When choosing the optimized configuration (green curve, $m_{\text{bat}} = 3.5\,\text{kg}$), both excess time and minimum state-of-charge with SAR payload increase significantly with respect to the original configuration (red curve).

3.1.3 Human Detection and Tracking Framework

Compared to our previous publications (Kümmerle et al., 2016) this paper demonstrates that the human detection and tracking framework also works in realistic search-and-rescue scenarios and onboard the limited sensing payload of fixed-wing UAVs such as *AtlantikSolar*. The framework can be split up into the four stages that are shown in Figure 5 and are explained in more detail in the following.

Preprocessing

The preprocessing stage takes the raw sensor outputs and finds synchronized images as well as the corresponding pose of the UAV. In addition, it converts the 14 bit-valued infrared image (where each pixel corresponds to a temperature between −40 °C and +550 °C) to an 8 bit image to guarantee compatibility with standard computer vision algorithms.

[2]First, batteries can only be added in 290 g packs due to the six-cell (6S) series configuration and second, y-axis mass-symmetry has to be retained.

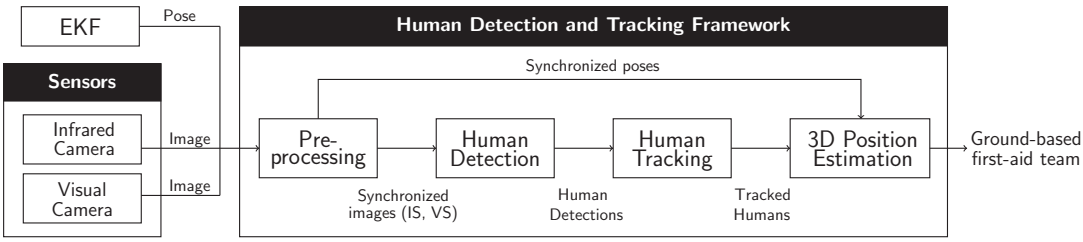

Figure 5: Overview over the complete human detection and tracking system with data flow.

Human Detection

The human detection pipeline depicted in Figure 6 first detects all potential humans in the infrared spectrum (IS) and then uses additional visual spectrum (VS) information to reduce the false positive rate.

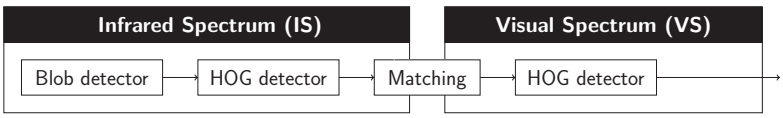

Figure 6: Overview over the complete human detection pipeline.

1. **Blob detector (IS):** The blob detector (Hinz, 2005; Bradski et al., 2000; Suzuki et al., 1985) aims to reliably detect *all potential* humans in the infrared image. It is therefore configured for a low miss rate rather than a low false positive rate. As humans often appear as textureless blobs in the infrared spectrum detecting these blobs results in a very high human detection rate. The brightness difference between each detected blob and its surroundings is computed. Higher brightness difference means higher contrast and therefore higher saliency. The maximum number of detected blobs per image is limited to make the framework's runtime more predictable.
2. **HOG detector (IS):** In the second stage, a Histogram of Oriented Gradients (HOG) detector (Dalal & Triggs, 2005) is applied to all first-stage candidates. At every blob location, candidate bounding boxes with all relevant human sizes are created (Figure 7a,b) and classified by the HOG detector. The scores of candidates centered at the same location are compared to each other. If the highest score is negative then the location is classified as non-human, if it is positive the corresponding candidate is forwarded to the next stage. To train the HOG detector infrared images from four infrared datasets (Wu, Fuller, Theriault, & Betke, 2014; Davis & Keck, 2005; Davis & Sharma, 2007; Vempati, Agamennoni, Stastny, & Siegwart, 2015) with a total of 2400 positive and 24000 negative human samples[3] were used. For the training of the linear Support Vector Machine (SVM) the auto training function in (Bradski et al., 2000) is employed. It automatically finds good training parameters by optimizing over the cross-validation (10-fold) error. For testing performance, a set of 400 positive and 4000 negative samples[3] was used.
3. **Matching (IS to VS):** The matching step maps the detections from the infrared- to the visual spectrum. This is done by undistorting and vertically rectifying both images. The correspondence search is then trivially solved by taking the pixels with same coordinates.
4. **HOG detector (VS):** The last stage consists of a visual spectrum HOG detector. The candidate bounding boxes from the infrared HOG detector are already converted to the visual spectrum image by the matching stage, but due to inaccurate camera calibration the matching is not perfect in practice. Therefore, not only the direct matching candidate but also further candidates sampled in its neighborhood (Figure 7c) are classified by the HOG detector. The focus can be set on precision or runtime by using small or large distances between the candidates respectively. The final decision for a region is positive (human detected) if the classification output y for any

[3]Including left-right reflections.

candidate in that region exceeds a user-specified threshold. Only if none of the candidates in a region is classified as human the final decision is negative (no human detected). The training of the HOG detector is done as described in step two. A new training dataset consisting of 3000 humans[3] and 25000 non-human patches was collected from various publicly available images.

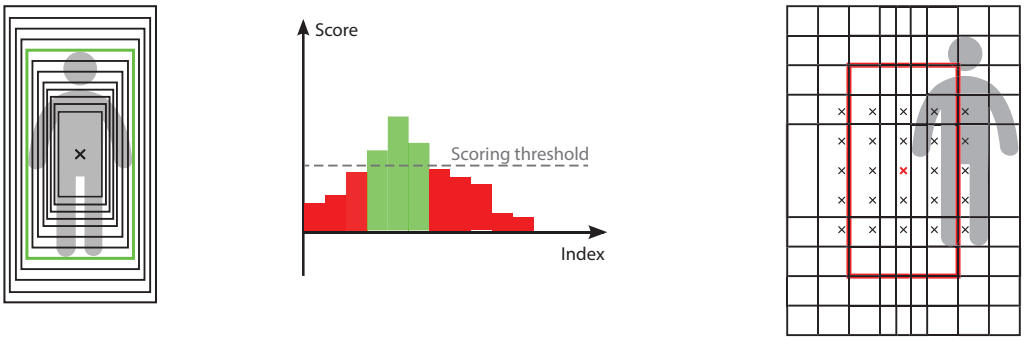

(a) Candidate bounding boxes of different size centered on a blob detection (marked by a cross) on a human. The candidate with highest classification score is highlighted with green color.

(b) Classification score for different candidates, which are classified as human or non-human if they exceed (green) or stay below (red) the scoring threshold respectively.

(c) Sampling structure in the visual spectrum image for a matched candidate (red rectangle). The centers of the rectangles are marked by crosses. The positions of the samples result from horizontal and vertical shifting of the red rectangle by fixed distances.

Figure 7: Sampling around human detections in the infrared- and visual spectrum.

Human Tracking and 3D Position Estimation

The *Human Tracking* stage compares the human detections in the current frame to the humans already tracked in previous frames. Those not tracked yet are registered for tracking as of the next frame. For tracking several humans at the same time a Kalman filter with a Binarized Normed Gradient (BING) objectness model (Cheng, Zhang, Lin, & Torr, 2014) is employed. Every detection creates a filter that tracks the human observations from frame to frame. This functionality is described in detail in our publication (Kümmerle et al., 2016). The 2D location of all tracked humans in the infrared image is then provided to the *3D Position Estimation* stage, which, based on pairs of UAV pose and 2D human position calculates the respective 3D positions in world coordinates and forwards them to the ground-based rescue teams.

3.2 Autonomous Launch and Landing

To expand applicability of small solar-powered UAV platforms, launch and landing processes should be automated such that experienced pilots are not needed to safely operate the aircraft. This includes operation in high winds and atmospheric turbulence, where an autopilot with proper environmental sensing has potential for improved safety during critical near-terrain flight. The fidelity of landing and take-off automation are especially pertinent for small, solar-powered UAVs due to their dynamics (e.g. tight bank angle and sink rate constraints) as well as their relatively low structural robustness. Delicate handling of all procedures and contingencies is required. The automatic launch and landing framework described in this paper is based on the implementation in the *Pixhawk* autopilot (Meier et al., 2015). In this section, we will describe our extensions to this code base, i.e. additional launch criteria for hand-launch, landing *go-around* logic, spiral-down paths, final approach entrance conditions/contingencies, and integration of LIDAR-based terrain sensing.[4]

[4]Further details on the high-level control structure are available at http://pixhawk.org/.

Table 2: Parameters for constructing the landing path.

Parameter	Description
h_{home}	Home altitude (absolute)
$(lat, lon)_{\text{land}}$	Landing point
h_{des}	Start descent altitude (AGL)
h_{fl}	Flare altitude (AGL)
h_{virt}	Logarithmic virtual height*
h_{app}	Final approach entrance altitude (AGL)
χ_{app}	Final approach bearing
γ_{slope}	Landing slope angle
R_{loit}	Loiter down radius
dir_{loit}	Loiter down direction (1=clockwise, -1=counter-clockwise)

*see further details about the construction of the flare curve on the *Pixhawk* webpage: https://pixhawk.org/users/fixedwing_autoland.

3.2.1 Automatic Hand-Launch

The most critical portion of an automated hand-launch is the detection phase. Hand-launchers may have a range of throwing strengths that, on the one hand, could trigger a *false positive* detection that enables the motors and potentially results in injuries to the operator, and on the other hand a *false negative* when the aircraft has in fact left the hand of the launcher could result in damage to the platform itself. The existing automated launch detection module in the *Pixhawk* autopilot waits for a positive x-body axis acceleration a_x (in line with the motor) crossing a threshold a_x^{thres} over the time interval T_{launch}. Obviously, a_x^{thres} is an important parameter that depends on both the strength of the given operator's typical throw and, for low-speed solar-powered UAVs, whether an excessive launch acceleration is even needed. In the case of *AtlantikSolar*, nominal operating speed is $9.6\,{}^{\text{m}}/_{\text{s}}$, a speed achievable with minimal throwing strength given a running start. To account for such factors, a parallel launch detection criteria that monitors whether the airspeed v_{air} exceeds a given threshold v_{thres} over T_{launch} was introduced. In high wind launches this criteria should be carefully chosen and combined with the a_x-threshold to avoid unexpected motor activity during the running start.

Once a launch is detected, high-level guidance commands a climbout with maximum allowed throttle set to full (cruise often has a maximum throttle value set much lower) and level roll until such point the aircraft reaches sufficient altitude $h_{\phi_{\text{lim}}^{\text{TO}}}$ to allow banking without fear of tip-strike, or loss of lift during the initial take-off. A small positive pitch limit $\theta_{\text{lim}}^{\text{TO}}$ also bounds pitch commands to remain in an ascent regime. Throttle resumes standard cruise controlled commands once some vertical tolerance (climbout difference, $\Delta h_{\text{climbout}}$) of the first waypoint `wp0` is reached. A detailed take-off process diagram is shown in Figure 8.

3.2.2 Automatic Landing

Automatic Landing Setup

The landing sequence can be constructed with the parameters in Table 2. Figure 9 shows an exemplary race-car pattern with spiral-down ends that guarantees that the safety pilot is always in visual-line-of-sight of the aircraft. The segments of the path can also be seen as Dubins Aircraft path segments. Note that, given an unobstructed approach vector, a simple linear descent could also be executed; though, including spiral-down logic extends practically to more general situations where trees or other structures may limit approach directions, or airspace allowing low-altitude flight may be more constrained.

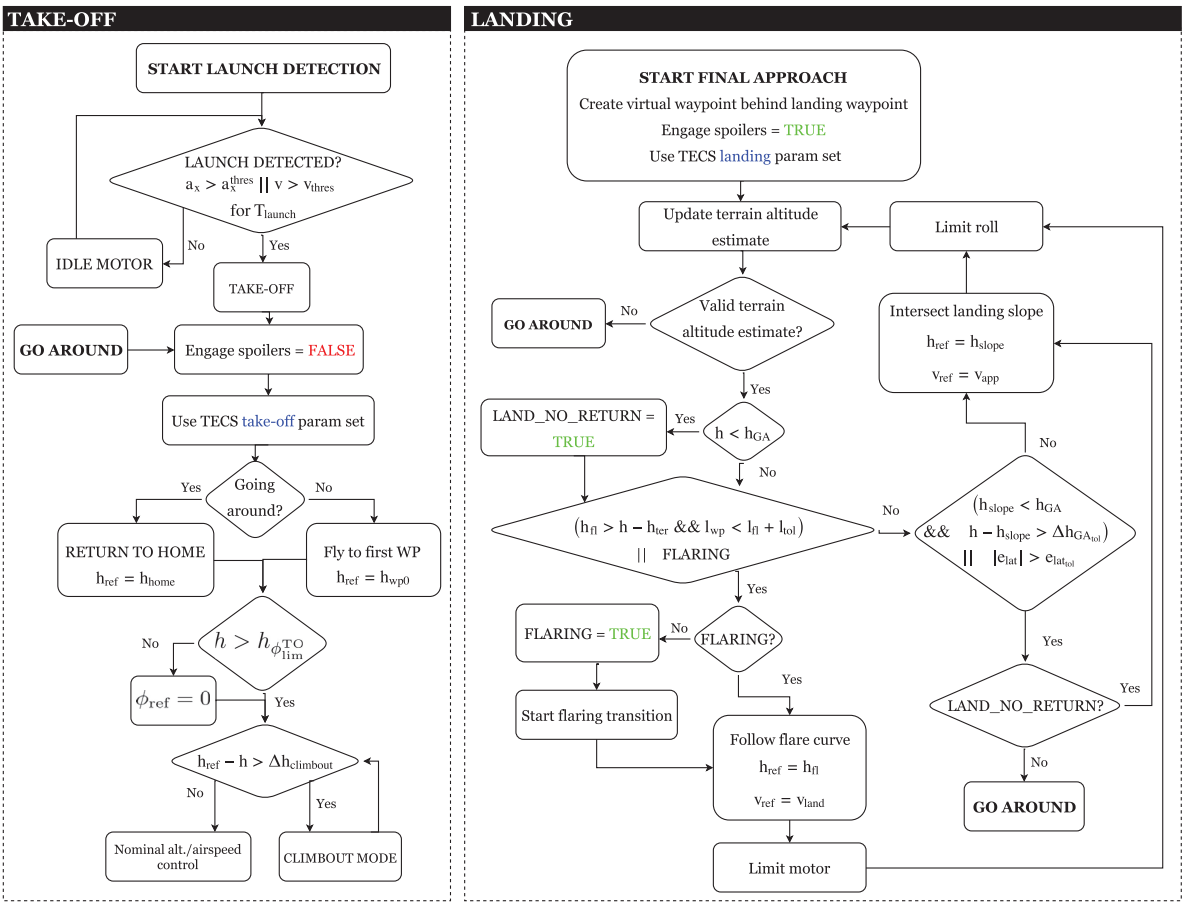

Figure 8: Automatic take-off and landing process diagrams.

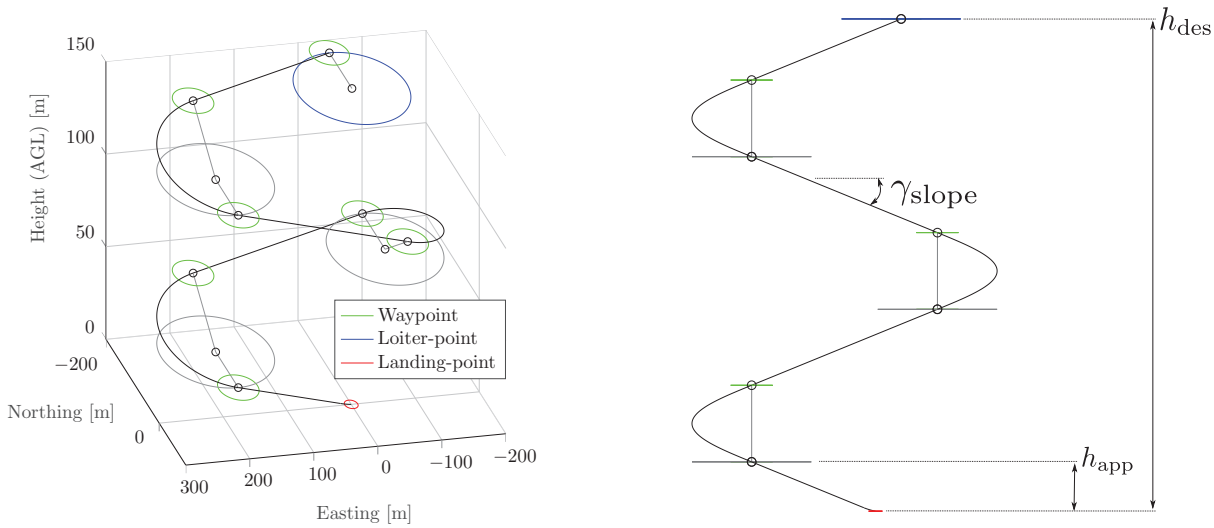

(a) (b)

Figure 9: Visualization of an exemplary "race-car" landing descent sequence with (a) a three-dimensional view and (b) a side view with relative altitude definitions.

High-level Descent and Landing Guidance

The high-level descent and landing guidance uses time-independent path following (as opposed to trajectory following) to improve robustness against atmospheric disturbances (static winds, gusts, thermals, ground-effects) that influence especially slow solar-powered UAVs during landing. As in Figure 3, we employ a $\mathcal{L}_1$-based approach for lateral-directional path following and TECS for altitude and airspeed control. As glide paths for such a high-lift aircraft must be kept shallow to ensure tracking even in potential vertical updrafts, we may assume this decoupled structure sufficient for tracking these three-dimensional curves. A necessary modification to the standard high-level logic is however position-based spiral-down logic: Inside a spiral-down segment, the TECS altitude reference is set based on the angular travel from the starting waypoint at the top of the spiral until reaching the next waypoint, according to the spiral height equation $h_{\text{slope}} = h_{\text{last-wp}} - R_{\text{loit}}\Delta\xi/\pi$. Note that we only consider $180\,^\circ$ travel before entering the next straight segment of the descent.

The spiral-down logic switches to the next straight segment when the aircraft enters the current waypoint's acceptance radius. A simple constant vertical tolerance is used to ensure proper vertical path tracking. If the acceptance radius or vertical tolerance is missed (e.g. due to a wind gust), *try again* contingency logic is implemented such that the aircraft circles around and attempts to rejoin the path, entering the acceptance volume on a second try (Figure 10a). The altitude reference for any angular travel beyond the current waypoint (e.g. in our case, anything greater than $180\,^\circ$) is set to that of the last waypoint, $h_{\text{ref}} = h_{\text{last-wp}}$. Acceptance radius based switching also induces a corner case when entering the spiral for the first time, when the aircraft is still in front of the waypoint. Here, setting the altitude reference to the descent slope is necessary to avoid step changes in the reference which may disturb the current tracking of the slope (see the small blue region in Figure 10a).

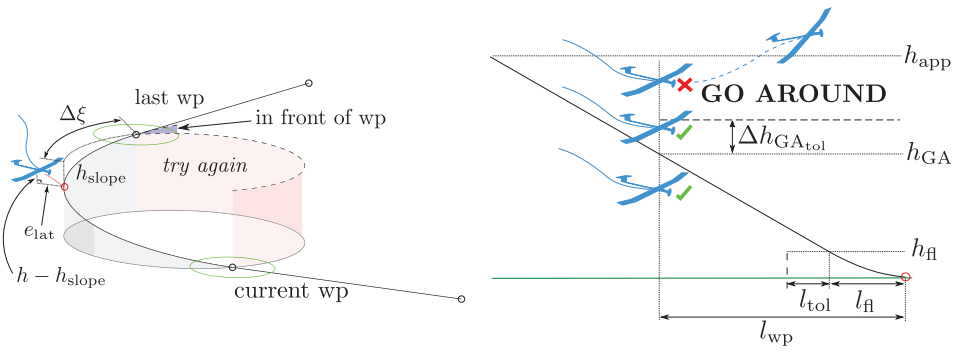

(a) Spiral-down logic.

(b) Final approach definitions and *go-around* logic.

Figure 10: Descent and landing logic with contingencies.

Tracking an inertially defined descent path with a high-lift vehicle is challenging when thermal updrafts or turbulence are present. In addition, altitude controllers for solar-powered UAVs such as *AtlantikSolar* are usually only loosely tuned to allow a smooth throttle behavior and thus efficient cruise. To guarantee good vertical tracking during landing a separate set of TECS tuning parameters was thus implemented. For altitude loop aggressivity the height rate feed-forward gain $k_{\dot{h}_{\text{ff}}}$ was increased significantly, the throttle time constant τ_T was set approximately an order of magnitude lower, and the integral gain k_{I} was increased in the landing parameter set to quicken altitude offset regulation. Further, engaging spoilers during descent was required to increase the glide-slope angle. Despite the more aggressive tuning, strong gusts and updrafts can still cause excessive tracking offsets. Especially during the final approach, it is important to include contingencies if the aircraft is not heading for the desired landing point. We, therefore, implemented *go-around* logic during the final approach phase in the event either lateral or vertical track error tolerances, $e_{\text{lat}_{\text{tol}}}$ and Δh_{GA}, respectively, are exceeded. The process diagrams shown in Figure 8 detail the *go-around* logic and Figure 10b shows the final approach definitions pictorially. Further, roll limitations are set as a

function of altitude (above ground) in all conditions to avoid tip-strike. Note symbols v_{app} and v_{land} are the airspeed setpoints (typically the same value) for the landing approach and flare, respectively.

Altimetry

GPS, without the use of differential GPS (D-GPS), does not provide the level of altitude accuracy necessary for autonomous precision landing (including altitude-dependent flaring and motor shutoff) with small solar-powered UAVs. Another solution is the use of laser-altimetry. We employ a compact and lightweight LIDAR altimeter[5] with up to 40 m range that is integrated into *AtlantikSolar*'s fuselage pointing nadir. The LIDAR, which consumes approximately 1.5 W, is switched-off during nominal operation and only turned-on when beginning the landing procedure. The sensor's reliable range starts at ca. 25 m. A typical final approach altitude for *AtlantikSolar* is $h_{\text{app}} = 18\,\text{m}$. We therefore limit use of the LIDAR measurements for terrain estimation to the final approach, where roll and pitch angles typically remain small enough to approximate altitude above ground. A simple outlier rejection is used to handle erroneous sensor measurements, and a timeout on previous "valid" measurements is used for abort criteria. Specifically, if the sensor either gives no measurement (sensor fault) or erroneous measurements (a deviation of greater than some threshold from the last value) for some time, the final approach is aborted and a *go-around* is initiated.

3.3 Thermal Updraft Exploitation Applied to Solar-Powered UAVs

The exploitation of thermal updrafts is a well-known method to extend aircraft flight endurance and range. The origins of manned thermal soaring date back to the 1930s (U.S. Federal Aviation Administration, 2013). Autonomous soaring was only suggested by Wharington (1998) much more recently. Wharington adapted the simple heuristic rules for piloted soaring by Reichmann (1988) with gain-scheduling approaches trained through reinforcement learning. He was also the first to suggest the combination of autonomous thermal soaring with solar-powered flight. Allen (2006, 2007) formulated a comprehensive altitude-dependant bell-shaped thermal model, developed the well-known *centroid method* and provided first extensive flight results for autonomous thermal soaring. Edwards and Silverberg (2010) and Edwards (2015) refined Allen's batch method with the *simultaneous iteration* method. The computations were performed on a ground-based computer, but up to five hour long soaring flights were shown in a cross-country soaring challenge against manual RC-glider pilots. Recursive filtering approaches that can handle the non-linearities of thermal estimation were introduced through a lightweight Kalman filtering method (Hazard, 2010) and a more general particle filter approach (Bencatel, 2010). All these approaches share certain observability issues, i.e. they need to estimate several quantities (thermal strength, radius, latitude and longitude) from only a single measurement: The local updraft speed w. Authors have therefore suggested to integrate updraft field gradient information. Fonseka (2007) actively used the roll acceleration. Li (2010) employed the heading angle offset caused when entering the thermal to determine the initial loitering direction, but not to improve the estimation process itself. Bower, Flanzer, Naiman, and Saripalli (2010) introduced the roll moment measurement into the estimation problem, but do not answer how to measure the roll moment in flight. None of this work has deployed and analyzed a roll moment based estimation scheme in actual flight tests.

The following sections implement a novel lightweight, Kalman filtering based approach for autonomous thermal updraft tracking into the well-known *Pixhawk* autopilot. Our main research contributions are: First, we present an end-to-end method to integrate the thermal-induced roll moment. This additional measurement improves the problem's observability and is especially suitable for solar-powered UAVs due to their usually high wingspan. We provide the first analysis of the approach's performance based on actual flight results. Second, we extend the local updraft speed measurement with a term that considers the load factor n_z and thus the increased sink rate during turns. Third, 20 years after Wharington (1998) first highlighted the benefits of combining autonomous soaring techniques with solar-powered flight — especially to enable *perpetual-flight* applications — this paper marks the first field deployment of such kind.

[5] LIDAR-Lite v2 from Pulsedlight, https://www.pulsedlight3d.com/products/lidar-lite-v2-blue-label.html.

3.3.1 An Extended Thermal Model

The updraft model is based on the commonly used mathematical representation of a single thermal updraft found in (Wharington, 1998; Hazard, 2010). A thermal updraft at time step k is represented through its state vector

$$X_k = [W \ R \ r_\mathrm{n} \ r_\mathrm{e}]^T \ , \tag{1}$$

where W represents the thermal's maximum or core updraft velocity, R the thermal radius, and r_n and r_e are the position of the thermal relative to the aircraft in north- and east-direction (Figure 11). The thermal dynamics presented by Hazard (2010) have been extended to account for the effects of wind: The thermal is assumed to be a *thermal tube* that starts at a fixed location on the ground and drifts along with the wind as altitude increases. Combining aircraft movement and wind drift, the thermal dynamics in the aircraft-relative coordinate system are

$$X_k = f(X_{k-1}) = X_{k-1} + \begin{bmatrix} 0 \\ 0 \\ -v_\mathrm{n} + v_\mathrm{n}^\mathrm{wind} \cdot \frac{v_\mathrm{d}}{v_\mathrm{d,th}} \\ -v_\mathrm{e} + v_\mathrm{e}^\mathrm{wind} \cdot \frac{v_\mathrm{d}}{v_\mathrm{d,th}} \end{bmatrix} \cdot \Delta t \ . \tag{2}$$

Here, v_n, v_e and v_d represent the components of the aircraft ground velocity vector and $v_\mathrm{n}^\mathrm{wind}$ and $v_\mathrm{e}^\mathrm{wind}$ represent the wind vector components (direction where the air is flowing to) in the inertial north-east-down (NED) reference frame. The variable $v_\mathrm{d,th}$ corresponds to the vertical velocity of the whole thermal structure (so usually $v_\mathrm{d,th} < 0$ because the thermal structure ascends). It is implemented as a parameter because its magnitude is not necessarily equal to the updraft core velocity W. Finally, the updraft speed distribution within the thermal is assumed to be normally distributed:

$$w(r_\mathrm{n}, r_\mathrm{e}) = W \cdot e^{\left(-\frac{r_\mathrm{n}^2 + r_\mathrm{e}^2}{R^2}\right)} \ . \tag{3}$$

The local updraft speed w and the resulting updraft gradient dw/dr are shown versus the radius $r = \sqrt{r_\mathrm{n}^2 + r_\mathrm{e}^2}$ in Figure 12. Note that we do not include a sink term in the model because calculating the average sink speed requires solving the vertical air mass conservation equation, which in turn requires knowledge of the updraft distribution over the full area and not only that of a single thermal.

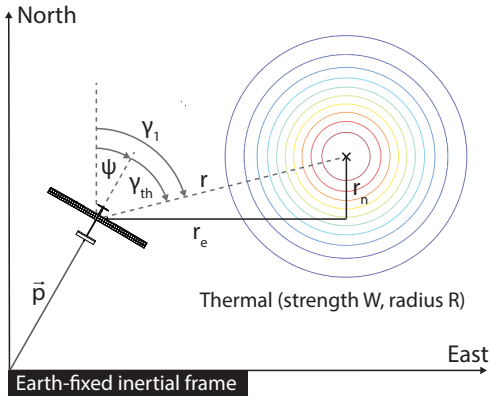

Figure 11: Coordinate systems and horizontal thermal geometry.

The non-uniform updraft velocity distribution in a common thermal induces a moment around the aircraft pitch- and roll-axes. The influence of the latter is easier to determine because of, first, the high aspect ratio wings and thus larger lever arm found on typical soaring aircraft and, second, the fact that the pitch dynamics are frequently manipulated by the autopilot to control the airspeed. This paper therefore focuses on the roll moment induced by a thermal. It is retrieved by integrating over the non-uniform pressure

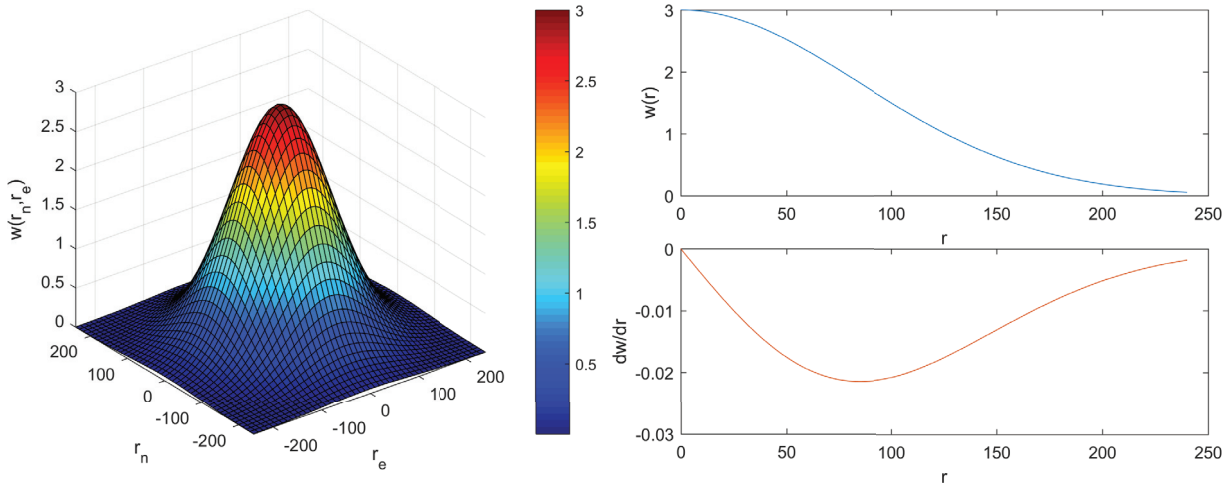

Figure 12: Model representing a thermal updraft with $W = 3\,\mathrm{m/s}$ and $R = 120\,\mathrm{m}$. Thermals with these characteristics are typical at low altitudes above ground and European latitudes in summer. Left: Isometric view of the updraft velocity $w(r_\mathrm{n}, r_\mathrm{e})$. Right: Local updraft speed $w(r)$ and updraft gradient dw/dr.

distribution change $\Delta p(b)$ along the span-wise position b due to the thermal updraft

$$L_\mathrm{th} = -\int_{-B/2}^{B/2} \Delta p(b) c_\mathrm{wing} b \, db \tag{4}$$

$$= -\frac{1}{2}\frac{dc_\mathrm{L}}{d\alpha}\frac{dw}{db}(W, R, r_\mathrm{n}, r_\mathrm{e}) \rho v_\mathrm{air} c_\mathrm{wing} \int_{-B/2}^{B/2} b^2 db \tag{5}$$

$$= -\frac{1}{24}\frac{dc_\mathrm{L}}{d\alpha}\frac{dw}{db}(W, R, r_\mathrm{n}, r_\mathrm{e}) \rho v_\mathrm{air} c_\mathrm{wing} B^3 \ . \tag{6}$$

Here, B is the wing span and c_wing is the wing chord length. To expand Eq. (4) we use the relationship $\Delta p(b) = \frac{1}{2}\rho v_\mathrm{air}^2 \Delta c_\mathrm{L}(b)$ where ρ is the air density and Δc_L is the local lift coefficient change. The dependence of the lift coefficient on the angle of attack α can (in normal cruise flight and thus at sufficiently small α) be linearized such that $\Delta c_\mathrm{L}(b) = \frac{dc_\mathrm{L}}{d\alpha} \cdot \Delta\alpha(b)$. The wing or aircraft lift-curve slope $\frac{dc_\mathrm{L}}{d\alpha}$ can be determined from standard aerodynamics literature (Pamadi, 2004; Raymer, 2006) or through aerodynamic modeling tools[6]. To retrieve $\Delta\alpha(b)$, we (again for small α) note that the updraft speed $\Delta w(b)$ causes a change in local α according to Figure 13a. If $B \ll R$, which is always the case for thermals that are large enough to be exploitable by a specific UAV, then $\Delta w(b) \ll v_\mathrm{air}$ (for the *AtlantikSolar* UAV in Figure 12, $\max(\Delta w) = 0.077\,\mathrm{m/s}$ and $v_\mathrm{air} \approx 9\,\mathrm{m/s}$) and thus $\tan(\Delta\alpha(b)) \approx \Delta\alpha(b) = {}^{\Delta w(b)}/_{v_\mathrm{air}}$. Furthermore, it is then permissible to linearize the updraft distribution such that $\Delta w(b) = \frac{dw}{db} \cdot b$. Substituting these relationships, rearranging and then executing the integration over the wing span (see Figure 13b) yields Eq. (6).

All terms of Eq. (6) except for the updraft distribution gradient $\frac{dw}{db}$ are independent of the thermal state. As shown in Figures 11 and 13, $\frac{dw}{db}$ depends both on the thermal state and aircraft attitude (yaw angle ψ and roll angle ϕ). It is retrieved by calculating the gradient of Eq. (3) along the wing span:

$$\frac{dw}{db} = -\frac{dw}{dr}\cos(\phi)\sin(\gamma_\mathrm{th}) = \frac{2W}{R^2} e^{\left(-\frac{r_\mathrm{e}^2 + r_\mathrm{n}^2}{R^2}\right)} \cos(\phi)\left[\cos(\psi) r_\mathrm{e} - \sin(\psi) r_\mathrm{n}\right] \ . \tag{7}$$

Here, $\gamma_\mathrm{th} = \gamma_1 - \psi = \arcsin(\frac{r_\mathrm{e}}{r}) - \psi$ is the horizontal angle between the aircraft longitudinal axis and the thermal center. Using the vector cross product between aircraft yaw and thermal direction we can find the alternative expression

$$\sin(\gamma_\mathrm{th}) = \tfrac{1}{r}(\cos(\psi) r_\mathrm{e} - \sin(\psi) r_\mathrm{n}) \ , \tag{8}$$

[6] E.g. *XFLR5*, available at http://www.xflr5.com.

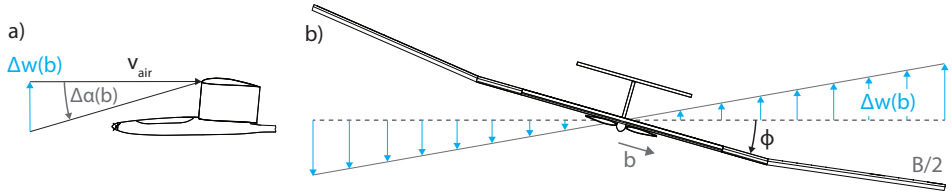

Figure 13: Left: Angle of attack change $\Delta\alpha(b)$ due to the thermal-induced local updraft velocity $\Delta w(b)$. Right: Rear-view of the aircraft and the assumed linear distribution of $\Delta w(b)$ along the wing. The aircraft is in a clockwise turn around a right-hand sided thermal. In this example, L_{th} is thus negative.

which together with the evaluated $\frac{dw}{dr}$ results in Eq. (7). Note that the trigonometric terms are not a function of the thermal state anymore, thus avoiding the otherwise cumbersome differentiation in the implementation of the EKF update (Section 3.3.3). Combining Eqs. (6) and (7), we retrieve the final form of the expected induced roll moment in a thermal

$$L_{\text{th}} = -\frac{1}{12}\frac{dc_{\text{L}}}{d\alpha}\rho v_{\text{air}} c_{\text{wing}} B^3 \frac{W}{R^2} e^{\left(-\frac{r_{\text{n}}^2+r_{\text{e}}^2}{R^2}\right)} \cos(\phi)\left[\cos(\psi) r_{\text{e}} - \sin(\psi) r_{\text{n}}\right] \ . \tag{9}$$

3.3.2 Thermal Sensing Onboard the Aircraft

Local Updraft Speed Sensing

The estimation of the local updraft speed w is based on the total specific energy $\varepsilon = {}^{E}/_{mg} = h + {}^{v^2}/_{2g}$. Defining the vertical speed $v_z = -v_{\text{d}}$ as positive upward for simplicity, the aircraft change in total specific energy is

$$\dot{\varepsilon} = v_z + \frac{\dot{v}v}{g} \ . \tag{10}$$

In contrast to existing literature (e.g. Hazard, 2010), we are not just interested in determining the location of the thermal center (where $\dot{\varepsilon}$ has a maximum), but we require a physically correct estimate of w and thus W to later infer the induced roll moment L_{th}. Therefore, this paper extends existing approaches by adding a correction term for the expected sink rate to retrieve the *measured*[7] local updraft speed

$$\hat{w} = \hat{\dot{\varepsilon}} - v_z(v_{\text{air}}, n_z, u_{\text{thr}}) \ . \tag{11}$$

Here, $v_z(v_{\text{air}}, n_z, u_{\text{thr}})$ represents the *expected* aircraft vertical velocity caused by drag minus propulsion system contributions. It is visualized in Figure 14. To identify the v_{air}-dependency, i.e. the well-known aircraft *sink curve*, the sink rate is measured in calm air at three different airspeeds (the default parameters $v_{\text{air}}^{\text{min}}$,$v_{\text{air}}^{\text{nom}}$ and $v_{\text{air}}^{\text{max}}$ of the *Pixhawk* autopilot) and then interpolated to retrieve $v_z(v_{\text{air}})$ in later flights. The load factor n_z correction accounts for higher drag in maneuvering- or turning-flight (such as during thermal centering) where the roll angle ϕ causes a load factor $n_z = 1/\cos(\phi)$ under a coordinated turn assumption. The expected v_z at airspeed v_{air} and load factor n_z is derived in Appendix A and results in

$$v_z(v_{\text{air}}, n_z) = n_z^{3/2} v_z\left(\frac{v_{\text{air}}}{\sqrt{n_z}}\right) \ , \tag{12}$$

where $v_z(v_{\text{air}})$ is the vertical speed from the sink curve evaluated at $n_z = 1$. To correct for throttle effects, we again use standard *Pixhawk* parameters to perform a simple piecewise linear correction of v_z. The parameters are specified such that $v_z(v_{\text{air}}^{\text{nom}}) = 0$ at $u_{\text{thr}}^{\text{cruise}}$ and $v_z(v_{\text{air}}^{\text{nom}}) = v_{\text{climb}}^{\text{max}}$ at $u_{\text{thr}}^{\text{max}}$. Our approach

[7] This paper always applies the $\wedge$-superscript to highlight that the quantity is — in that specific context — a measurement.

crudely approximates the strong non-linearities of the throttle to climb-rate response visible in Figure 14. Depending on u_{thr}, we apply one of the following expressions to retrieve the final v_z:

$$v_z(v_{\text{air}}, n_z, u_{\text{thr}}) = \begin{cases} v_z(v_{\text{air}}, n_z) - v_z(v_{\text{air}}^{\text{nom}}) \cdot \frac{u_{\text{thr}}}{u_{\text{thr}}^{\text{cruise}}} & \text{if } u_{\text{thr}} \leq u_{\text{thr}}^{\text{cruise}} \\ v_z(v_{\text{air}}, n_z, u_{\text{thr}}^{\text{cruise}}) + v_{\text{climb}}^{\max} \cdot \frac{u_{\text{thr}} - u_{\text{thr}}^{\text{cruise}}}{u_{\text{thr}}^{\max} - u_{\text{thr}}^{\text{cruise}}} & \text{if } u_{\text{thr}} > u_{\text{thr}}^{\text{cruise}} \end{cases} . \tag{13}$$

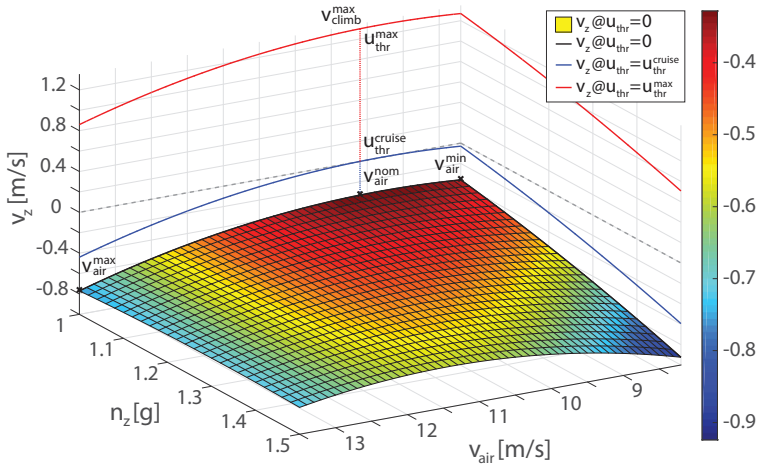

Figure 14: The proposed aircraft sink rate model, i.e. the expected vertical speed v_z as a function of airspeed v_{air}, load factor n_z and throttle u_{thr} visualized exemplarily for the *AtlantikSolar* UAV[8]. The colored surface shows v_z at $u_{\text{thr}} = 0\,\%$, the blue and red lines indicate v_z at $u_{\text{thr}}^{\text{nom}}$ and $u_{\text{thr}}^{\max}$ respectively.

Induced Roll Moment Sensing

The uncoupled, 1-DOF aircraft rotational dynamics (Raymer, 2006) about its longitudinal (x-) axis are mainly governed by the moment of inertia I_{xx}, the roll angle ϕ, the roll rate and acceleration p and $\dot{p}$, and the aileron deflection u_{ail}. When the roll motion is linearized about the trimmed flight condition ($\phi_0 = p_0 = 0$, $u_{\text{ail},0} = u_{\text{ail}}^{\text{trim}}$), then the roll rate dynamics can be expressed using the first-order differential equation

$$I_{\text{xx}}\dot{p} = L = C_{L_p} p + C_{L_{\text{ail}}} \left(u_{\text{ail}} - u_{\text{ail}}^{\text{trim}}\right) + L_{\text{th}} . \tag{14}$$

The left-hand side represents the roll moment L. The term $C_{L_p} p$ with $C_{L_p} = \delta L / \delta p < 0$ represents roll rate damping by the wing. The roll moment generated by the ailerons is considered by the control derivative $C_{L_{\text{ail}}} = \delta L / \delta u_{\text{ail}}$. Finally, L_{th} is the thermal-induced roll moment. The components of Eq. (14) can be used in different ways to gain information about the thermal. Human glider pilots usually focus on 'which wing tends to be lifted' (U.S. Federal Aviation Administration, 2013), i.e. they observe ϕ, p and $\dot{p}$ when entering the thermal and then turn to the opposite side of the induced roll motion. For unmanned gliders, Fonseka (2007) proposes a similar technique that focuses on the thermal-induced roll acceleration. However, atmospheric turbulence and the fact that common autopilot IMUs only provide $\hat{\dot{p}}$ via numeric derivation of p make the roll acceleration signal very noisy. Therefore, this paper focuses on the long-term effects of the induced roll moment. We assume a quasi-steady state, i.e. we neglect $\dot{p}$ in Eq. (14) — which is valid because our autopilot constantly stabilizes the roll angle such that the average $\dot{p}$ over our long measurement horizon tends to zero — and only use the roll rate and aileron command to measure

$$\hat{L}_{\text{th}} = -C_{L_p} p - C_{L_{\text{ail}}} \left(u_{\text{ail}} - u_{\text{ail}}^{\text{trim}}\right) . \tag{15}$$

[8]For *AtlantikSolar* we set $u_{\text{thr}}^{\text{nom}} = 40\,\%$, $u_{\text{thr}}^{\max} = 61\,\%$, $v_{\text{air}}^{\min} = 8.2\,\text{m/s}$, $v_{\text{air}}^{\text{nom}} = 9.6\,\text{m/s}$, $v_{\text{air}}^{\max} = 13.5\,\text{m/s}$ and measure corresponding sink rates at $u_{\text{thr}} = 0$ and $n_z = 1$ of $v_z(v_{\text{air}}^{\min}) = -0.36\,\text{m/s}$, $v_z(v_{\text{air}}^{\text{nom}}) = -0.33\,\text{m/s}$ and $v_z(v_{\text{air}}^{\max}) = -0.77\,\text{m/s}$.

Given that the roll rate p is small in coordinated turns at low pitch angle, our main indicator of L_{th} is u_{ail}. As an example, the thermals in Figures 11 and 13 both induce a roll moment to the left ($L_{\text{th}} < 0$) such that the aileron attempts to stabilize by rolling right via $(u_{\text{ail}} - u_{\text{ail}}^{\text{trim}}) > 0$. A central question to determine measurability is what *typical maximum* L_{th} and thus corrective u_{ail} we can expect in a thermal. By taking the derivative of Eq. (7) we find that the maximum updraft gradient and thus L_{th} occurs at $r = R/\sqrt{2}$. For the characteristic thermal in Figure 12 and our specific *AtlantikSolar* UAV configuration at $v_{\text{air}} = 9.6\,\text{m/s}$, the maximum roll moment $|L_{\text{th}}| = 2.95\,\text{Nm}$ occurs at $r = 84.9\,\text{m}$. To convert L_{th} to an expected corrective aileron action and vice-versa, the control and stability derivatives in Eq. (15) need to be determined. Linear system identification methods (Tischler & Remple, 2012) were applied to Eq. (14) for that purpose because they provide good estimation accuracy and allow to identify both C_{L_p} and $C_{L_{\text{ail}}}$ simultaneously and directly from flight data.[9] The flight test data needs to be carefully selected such that no thermal updrafts or winds bias the measurements, the aircraft shall be stabilized at the steady-state attitude (e.g. $\phi = \text{const}$ and $p \approx 0$) around which Eq. (14) is linearized, and low-amplitude step-inputs in u_{ail} shall be used because the aileron required to correct for the thermal-induced moment is also relatively small. Using this method, we retrieve $C_{L_p} = -52.1\,\text{Nm}$ and $C_{L_{\text{ail}}} = 51.1\,\text{Nm}$ for *AtlantikSolar* at $v_{\text{air}}^{\text{nom}}$. For $p \approx 0$ Eq. (15) shows that this specific UAV needs less than 6 % aileron deflection to counter the maximum roll moment induced by the thermal of Figure 12. Less aileron is needed in weaker thermals (smaller W), larger thermals (larger R), when not flying in the maximum updraft gradient region, or when the wings are not aligned with the updraft gradient.

Consequently, the estimation of $\hat{L}_{\text{th}}$ requires a very accurate model for $u_{\text{ail}}^{\text{trim}}$. Motivated by the proportionality $F_{\text{L}} \propto v_{\text{air}}^2$ we, first, consider the airspeed's effect on trim: The $u_{\text{ail}}^{\text{trim}}$ that leads to $\phi = \phi_{\text{ref}} = 0$ at $u_{\text{ail}} \approx 0$ is measured in autopilot-stabilized flight at $v_{\text{air}}^{\text{min}}$,$v_{\text{air}}^{\text{nom}}$ and $v_{\text{air}}^{\text{max}}$ such that $u_{\text{ail}}^{\text{trim}}(v_{\text{air}})$ can then be interpolated. Second, solar-powered UAVs usually employ a small wing dihedral (only 6 ° on *AtlantikSolar*) to achieve a solar module surface that is inclined equally to the sun everywhere. This increases the instability of the aircraft's spiral mode such that more corrective aileron needs to be applied inside a turn than in wings-level flight. As discussed in Section 4.4, $u_{\text{ail}}^{\text{trim}}(\phi)$ is best approximated by a sinusoidal with parameters a, b such that overall

$$u_{\text{ail}}^{\text{trim}}(v_{\text{air}}, \phi) = u_{\text{ail}}^{\text{trim}}(v_{\text{air}}) + a \cdot \sin(b \cdot \phi)\,. \tag{16}$$

3.3.3 Implementation

The thermal updraft tracking architecture consists of the three submodules shown in Figure 15. The *Measurements* module and its outputs $\hat{w}$ and $\hat{L}_{\text{th}}$ were described in Section 3.3.2. The *State Machine* module uses the local updraft speed $\hat{w}$ for thermal latch/unlatch decisions. The *Kalman Filter* module uses both measurements to estimate the current thermal state X_k. The whole framework is a plug-in module for the *Pixhawk* autopilot, i.e. it provides the current thermal state, suggested loitering direction and radius, and a data validity flag via a publish/subscribe-pattern (Meier et al., 2015) but does not change the higher-level guidance of the autopilot. Instead, *Pixhawk*'s $\mathcal{L}_1$ guidance decides independently whether and how a published thermal waypoint should be centered. Therefore, while executing standard missions (e.g. loitering, scanning or simple $A \rightarrow B$ missions) as usual, the framework can latch into a thermal to gain altitude and return to the previous mission e.g. once the maximum altitude is reached.

State Machine

The state machine has the important task to switch between the four modes `CRUISE/SEARCH`, `PRESOAR`, `SOAR` and `AVOIDSOAR` based on the boolean variable b_{latch}: If the local updraft velocity is larger than the user-specified limit w_{latch}, the altitude is below the maximum altitude h_{max} minus a margin, and the aircraft is not too far away from home, then $b_{\text{latch}} = \texttt{true}$. Note that the latch decision uses the strongly low-pass filtered $\hat{w}^{\text{ff}}$ to provide robustness against the significant noise in $\hat{w}$ (e.g. to avoid unnecessary and inefficient

[9] Note that many other methods exist: First, the aileron can be modeled as a deflected flat plate with $\alpha_{\text{plate}} = \delta_{\text{ail}}$, second, airfoil- and wing-simulation tools such as *XFLR5* (http://www.xflr5.com) allow a better (yet still first-order) consideration of the lift distribution along the wing, and third, empirical methods (Raymer, 2006; Pamadi, 2004) that include correction factors for the wing and aileron geometry exist. However, all these methods only provide crude initial approximations.

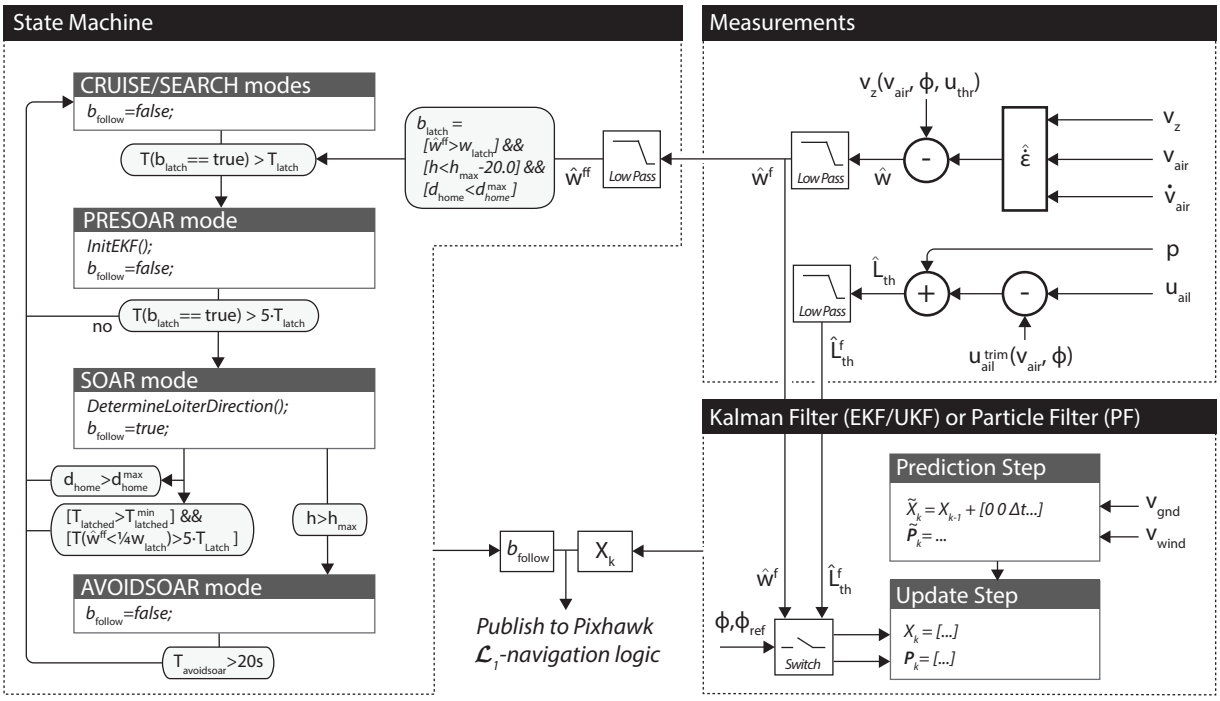

Figure 15: Architecture of the thermal updraft tracking framework.

turning maneuvers) while the Kalman filter uses the less filtered $\hat{w}^{\mathrm{f}}$ because it provides robustness against this (approximately Gaussian) noise by design. When no thermal updraft is present, the aircraft executes its standard mission in the `CRUISE/SEARCH` mode. As soon as $b_{\mathrm{latch}} = \texttt{true}$ for T_{latch} seconds[10], the aircraft enters `PRESOAR`. The *Kalman Filter* module is initialized and begins estimating the thermal state, however, the thermal is not actively tracked yet. The resulting advantage is that the state uncertainty has already decreased when the aircraft begins to actively track the thermal. This is the case in the `SOAR` mode (where $b_{\mathrm{follow}} = \texttt{true}$), which is engaged when $b_{\mathrm{latch}} = \texttt{true}$ for another $5 \cdot T_{\mathrm{latch}}$ seconds.[11] The aircraft determines the loiter direction based on the side at which the thermal is located and actively starts loitering around the thermal center at a user-specified radius R_{loiter}. If, while tracking the thermal, the aircraft exceeds the maximum allowed distance from home $d_{\mathrm{home}}^{\mathrm{max}}$, then the `CRUISE/SEARCH` mode is activated to fly the aircraft back into the operating area. If $\hat{w}^{\mathrm{ff}}$ is significantly below w_{latch} for a specified time, then `CRUISE/SEARCH` is also activated to either continue the previous mission or to search for other thermals. If the thermal updraft is however strong enough such that the altitude h reaches h_{max}, then the aircraft engages the `AVOIDSOAR` mode. This mode activates the spoilers and instantaneously directs the airplane back to its last nominal waypoint to get out of the thermal updraft. This functionality is of significant importance to avoid fly-aways.

Kalman Filter

The extended Kalman filter is a well-known recursive method to estimate the state of a system with a non-linear process- or measurement-function (Wan & Merwe, 2001). This section therefore only describes the specifics of applying an EKF for thermal updraft estimation. First, the filter is initialized once a thermal is detected (i.e. the `PRESOAR` mode is activated). We define the initial state

[10] Table 3 gives numeric values for the parameters w_{latch} and T_{latch} for medium-turbulence conditions during a mid-European Summer day for *AtlantikSolar*. The values are conservatively chosen and thus provide robustness to different atmospheric conditions. If truly optimal (e.g. faster) switching behavior is desired, then only w_{latch} and T_{latch} need to be adapted to the current atmospheric turbulence, the average thermal updraft diameter, and — to a lesser extent — the aircraft configuration.

[11] Note that the `SOAR` mode stays active for at least $T_{\mathrm{latched}}^{\mathrm{min}} = 0.8 \cdot {}^{2\pi R_{\mathrm{loiter}}}/_{v_{\mathrm{air}}^{\mathrm{nom}}}$ seconds because when it is activated the aircraft might have already passed the thermal center and then needs sufficient time to turn and fly back to the thermal core.

Table 3: Parameters chosen for the thermal updraft tracking approach.

Name	Value (Simulation)	Value (Flight)	Description
$v_{\mathrm{d,th}}$	-2.5 m/s	-2.5 m/s	Vertical (downwards) speed of the thermal bubble
R_{loiter}	80 m	80 m	Loiter radius around thermal
w_{latch}	0.6 m/s	0.6 m/s	Latch updraft velocity
T_{latch}	1.5 s	1.5 s	Latch time
Q_W	$(0.01\,\mathrm{m/s})^2$	$(0.01\,\mathrm{m/s})^2$	Process covariance for W
Q_R	$(0.25\,\mathrm{m})^2$	$(0.03\,\mathrm{m})^2$	Process covariance for R
Q_r	$(0.3\,\mathrm{m})^2$	$(0.2\,\mathrm{m})^2$	Process covariance for r_{n} and r_{e}
$R_{\hat{w}}$	$(0.2\,\mathrm{m/s})^2$	$(0.4\,\mathrm{m/s})^2$	Measurement covariance for $\hat{w}$
$R_{\hat{L}_{\mathrm{th}}}$	$(0.5\,\mathrm{Nm})^2$	$(3\,\mathrm{Nm})^2$	Measurement covariance for $\hat{L}_{\mathrm{th}}$
$P_{0,W}$	$(2\,\mathrm{m/s})^2$ [$(0.5\,\mathrm{m/s})^2$ for UKF]	$(2\,\mathrm{m/s})^2$	Initial covariance for W
$P_{0,R}$	$(80\,\mathrm{m})^2$ [$(20\,\mathrm{m})^2$ for UKF]	$(80\,\mathrm{m})^2$	Initial covariance for R
$P_{0,r}$	$(140\,\mathrm{m})^2$ [$(30\,\mathrm{m})^2$ for UKF]	$(100\,\mathrm{m})^2$	Initial covariance for r_{n} and r_{e}
$X_{0,W}$	1.5 m/s	2.5 m/s	Initial state component W
$X_{0,R}$	80 m	150 m	Initial state component R
D_0	30 m	30 m	Long. distance to thermal on initialization

$X_0 = [X_{0,W}, X_{0,R}, D_0\cos(\psi), D_0\sin(\psi)]^T$ such that the thermal is located at distance D_0 (set as a function of the average thermal radius) straight ahead of the aircraft. This guarantees good observability while the aircraft subsequently passes through the thermal. The initial covariance matrix $\boldsymbol{P}_0 = \mathrm{diag}(P_{0,W}, P_{0,R}, P_{0,r}, P_{0,r})$ is assigned large individual values (see Table 3) to represent the high initial thermal state uncertainty.

To propagate the thermal estimate in time, we use the process dynamics of Eq. (2) and assume that these are subject to additive, uncorrelated, time-independent zero-mean Gaussian noise u with process covariance matrix $\boldsymbol{Q} = \mathrm{diag}(Q_W, Q_R, Q_r, Q_r)$. The state-prediction[12] and covariance-prediction equations become

$$\tilde{X}_k = f(X_{k-1}) \tag{17}$$

$$\tilde{\boldsymbol{P}}_k = \boldsymbol{P}_{k-1} + \boldsymbol{Q} \ . \tag{18}$$

Note that due to the process dynamics $f(X_{k-1})$ given by Eq. (2), the process Jacobian $\boldsymbol{J}_f = df/dX = \mathbb{I}$ (where $\mathbb{I}$ is identity) and thus the covariance prediction has the simple form in Eq. (18). The EKF's update (or correction) step now corrects the predicted state $\tilde{X}_k$ and covariance $\tilde{\boldsymbol{P}}_{\boldsymbol{k}}$ using the difference between the actual measurement z and the expected measurement $h(\tilde{X}_k)$ as well as the Kalman gain $\boldsymbol{K}_k$:

$$X_k = \tilde{X}_k + \boldsymbol{K}_k\big[z_k - h(\tilde{X}_k)\big] \tag{19}$$

$$\boldsymbol{P}_k = \tilde{\boldsymbol{P}}_k - \boldsymbol{K}_k\boldsymbol{J}_h(\tilde{X}_k)\tilde{\boldsymbol{P}}_k \tag{20}$$

$$\boldsymbol{K}_k = \tilde{\boldsymbol{P}}_k\boldsymbol{J}_h^T(\tilde{X}_k)\big[\boldsymbol{J}_h(\tilde{X}_k)\tilde{\boldsymbol{P}}_k\boldsymbol{J}_h^T(\tilde{X}_k) + \boldsymbol{R}\big]^{-1} \ . \tag{21}$$

Here, we have used the definition of the on-board measurement vector

$$z = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \hat{w} \\ \hat{L}_{\mathrm{th}} \end{bmatrix} \ , \tag{22}$$

which we again assume is subject to the additive, uncorrelated, time-independent zero-mean Gaussian noise v with measurement covariance matrix $\boldsymbol{R} = \mathrm{diag}(R_{\hat{w}}, R_{\hat{L}_{\mathrm{th}}})$. The measurement expectation function $h(X)$ is defined by

$$h(X) = \begin{bmatrix} h_1(X) \\ h_2(X) \end{bmatrix} = \begin{bmatrix} w \\ L_{\mathrm{th}} \end{bmatrix} \ . \tag{23}$$

[12]Note that this paper always applies '~' to denote a forecasted quantity.

In both Eq. (22) and Eq. (23) we can use the results from the previous sections, i.e. Eqs. (11) and (15) and Eqs. (3) and (9) respectively. The measurement Jacobian $\boldsymbol{J}_h = dh/dX$ in Eqs. (20) and (21) can be obtained via straightforward derivation. The analytical form of $\boldsymbol{J}_h$ is given in Appendix B.

3.3.4 Observability Analysis

To analyse the information gain of the two measurements $\hat{w}$ and $\hat{L}_{\text{th}}$ observability analysis can be leveraged. Assessing the observability of non-linear systems in a formal manner is however not trivial (Tang et al., 2009) and beyond the scope of this paper. Instead, we use a more qualitative graphical approach. Note first that if only $\hat{w}$ is available, then we rely on only a single measurement to estimate four state variables. If $\hat{w}$ is measured at a fixed location relative to the thermal core, not all states are observable at the same time. However, measurements at different well-placed locations make the problem observable. To determine how observable each thermal state variable is from *one specific location* (instead of a full trajectory), we employ a graphical approach based on the linearized measurement update equations: Figure 16 plots the measurement Jacobian $\boldsymbol{J}_h$ used by the EKF. The Jacobian $\boldsymbol{J}_h$ is chosen because it determines whether the state covariance $\boldsymbol{P}_k$ in Eq. (20) stays bounded over time, of which observability is a sufficient condition (West & Harrison, 1997). More specifically, if $\boldsymbol{P}_k$ in Eq. (20) stays bounded over time, then for $\boldsymbol{Q} > 0$ the term $\Delta\boldsymbol{P}_k = \boldsymbol{P}_k - \tilde{\boldsymbol{P}}_k = -\boldsymbol{K}_k\boldsymbol{J}_h(\tilde{X}_k)\tilde{\boldsymbol{P}}_k < 0$ averaged over time. Plugging in Eq. (21) yields

$$\Delta\boldsymbol{P}_k = -\tilde{\boldsymbol{P}}_k\boldsymbol{J}_h^T\left[\boldsymbol{J}_h\tilde{\boldsymbol{P}}_k\boldsymbol{J}_h^T + \boldsymbol{R}\right]^{-1}\boldsymbol{J}_h\tilde{\boldsymbol{P}}_k\,. \tag{24}$$

Because $\boldsymbol{R} > 0$ the limits of this expression are $\lim_{|\boldsymbol{J}_h|\to 0}\Delta\boldsymbol{P}_k = 0$ and $\lim_{|\boldsymbol{J}_h|\to\infty}\Delta\boldsymbol{P}_k = -\tilde{\boldsymbol{P}}_k < 0$. To reduce the estimate's uncertainty, we therefore require a large magnitude in $\boldsymbol{J}_h$. A similar analysis for Eq. (19) shows that we also aim for $\boldsymbol{J}_h \gg 0$ to achieve a large $\boldsymbol{K}_k$ and thus a fast state correction (in the scalar case the maximum K_k does however not occur as $J_h \to \infty$ but at $J_h = \sqrt{R/P}$).

With this information, we can infer from Figure 16 that even when only using $\hat{w}$ the observability of the state component W is high in the thermal center where $dw/dW \approx 1$. Effectively, we are directly measuring the core strength W of the plotted typical thermal with $W = 3\,\text{m/s}$ and $R = 120\,\text{m}$. In other words, at $dw/dW = 1$ a $\Delta W = W - \tilde{W} = 0.4\,\text{m/s}$ or 16 % error in the state variable already causes a measurement deviation of $\Delta w = \hat{w} - w = 0.4\,\text{m/s}$ that is statistically significant (i.e. $|\Delta w| \geq |\sigma_{\hat{w}}|$ with $\sigma_{\hat{w}} = \sqrt{R_{\hat{w}}}$ from Table 3) when assuming a 1-σ or 68 % confidence level. The EKF can thus easily observe and correct this error according to Figure 16b: Given $\Delta w > 0$, the EKF increases the estimated W because $dw/dW > 0$ in the thermal. The observability of the thermal radius R is highest at $r = R$, where we have $dw/dR = 0.015$ (a $\Delta R = 26.7\,\text{m}$ or 22 % error causes $\Delta w = \sigma_{\hat{w}} = 0.4\,\text{m/s}$). The observability of the position r_n and r_e is of the same maximum magnitude at $r = \sqrt{R/2}$ (also see Figure 12), but the actual observability depends not only on the distance but also on the direction to the thermal core γ_{th}. As expected, for $r \gg R$ the $\boldsymbol{J}_h$ components decrease for all state variables and thus state estimation errors are not easily observable anymore.

However, even when measuring close to the thermal, paths that do not yield full observability using only w-measurements exist. Figure 16a) visualizes the following typical examples of paths that can be augmented using the L_{th}-measurement:

- Case A: *Circular path* $\mathtt{a}_1 \to \mathtt{a}_1$ *centered at the thermal core.* Assume the aircraft has perfectly encircled the thermal core $\mathtt{t}_1$ in Figure 16a. On this path, $\hat{w}$ is ambiguous in the sense that multiple combinations of thermal-strength and -size (e.g. large R / small W or small R / large W) can cause this w-measurement. The L_{th}-measurement removes this ambiguity because the mapping $(W, R) \to (w, L_{\text{th}})$ is unique.[13]
- Case B: *Straight-line path* $\mathtt{b}_1 \to \mathtt{b}_2$ *through the thermal core.* When the estimated thermal position is initialized at $\mathtt{t}_1$, the w-measurement does not yield any information about r_e because $dw/dr_\text{e} = 0$ along the path (Figure 16e), and thus the respective state covariance does not decrease. Using $\hat{L}_{\text{th}}$ allows the EKF to confirm that $r_\text{e} = 0$ and to reduce the state covariance.

[13] Note that Figure 16 subplots f)-j) are not applicable here, see the figure caption.

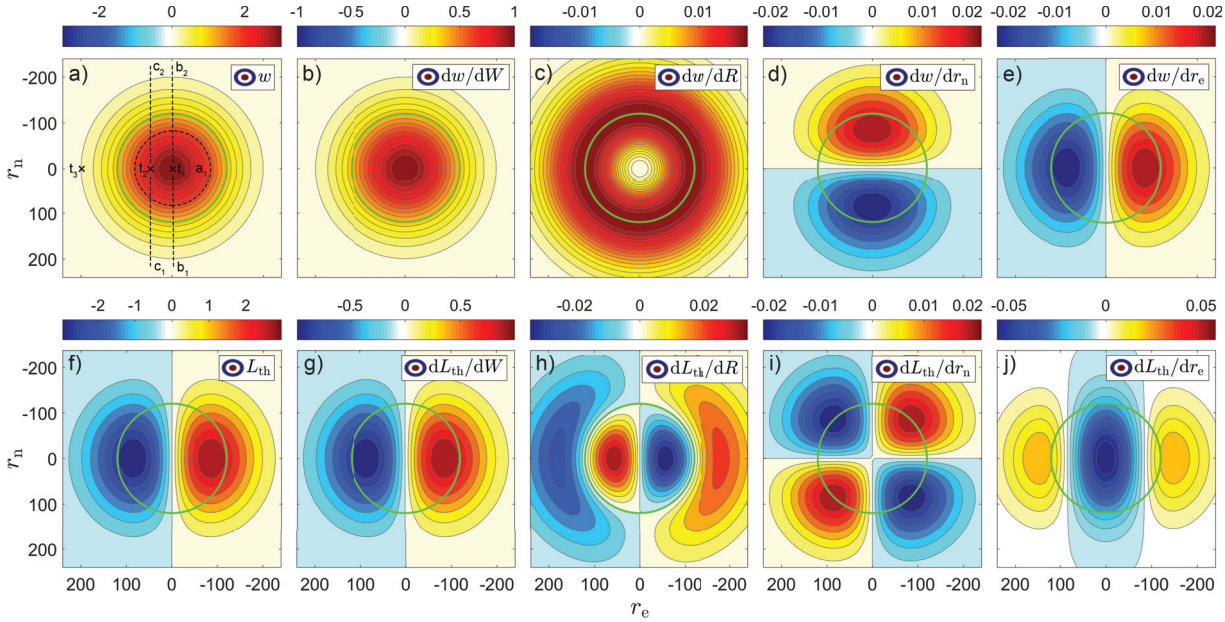

Figure 16: Thermal observability in the form of the measurement function $h(X)$ in subplots a) and f) and the measurement Jacobian dh/dX in subplots b-e) and g-j) over r_n and r_e. High magnitudes of those quantities mean good observability. The data is plotted for the thermal of Figure 12 with $W = 3\,\mathrm{m/s}$ and $R = 120\,\mathrm{m}$ (visualized by the green circle), the measurement noise $\boldsymbol{R}$ of Table 3 and the (already somewhat converged) state covariances $P_W = (0.5\,\mathrm{m/s})^2$, $P_R = (25\,\mathrm{m})^2$ and $P_r = (25\,\mathrm{m})^2$. The L_th-related subplots f-j) are only valid for an aircraft heading north, i.e. they assume $\psi = 0\,°$. The black dashed lines indicate hypothetical aircraft paths.

- Case C: *Straight-line path* $\mathtt{c}_1 \rightarrow \mathtt{c}_2$ *passing by the thermal core.* When the thermal is initialized at $\mathtt{t}_2$, $\hat{w}$ cannot provide a unique r_e-estimate (location $\mathtt{t}_2$ mirrored at $\mathtt{b}_1$-$\mathtt{b}_2$ is another solution). The EKF expects $L_\mathrm{th} = 0$ (Figure 16f) in this situation, measures $L_\mathrm{th} < 0$ because the thermal is to the right, therefore retrieves $\Delta L_\mathrm{th} = \hat{L}_\mathrm{th} - L_\mathrm{th} < 0$ and because $dL_\mathrm{th}/dr_\mathrm{e} < 0$ at that location (Figure 16j) the EKF correctly increases r_e and shifts the estimated thermal right.
- Case D: *Straight-line path* $\mathtt{c}_1 \rightarrow \mathtt{c}_2$ *with wrong initial thermal estimate.* Assume that, e.g. due to atmospheric turbulence, the thermal is wrongly initialized at $\mathtt{t}_3$. The fact that $\hat{L}_\mathrm{th}$ is negative while L_th is positive indicates that the estimated thermal lies on the wrong side of the aircraft. However, due to the linearization of $h(X)$ the EKF cannot benefit from $\hat{L}_\mathrm{th}$. Instead, with $\Delta L_\mathrm{th} < 0$ and $dL_\mathrm{th}/dr_\mathrm{e} > 0$ (Figure 16j), the EKF incorrectly shifts the thermal more to the left. This demonstrates the drawbacks of linearization-based filters such as the EKF and UKF. Fully non-linear filtering techniques that *can* return the correct solution are presented in Section 3.3.5.

Of course $\hat{L}_\mathrm{th}$ continuously adds information into the estimation process and not just in these four cases. When comparing $\hat{w}$ and $\hat{L}_\mathrm{th}$ over all regions of Figure 16 the absolute magnitude of their $\boldsymbol{J}_h$-components and thus their contributions to observability seem similar. However, this is solely a result of our assumptions: While our observability analysis is *global* over the shown r_n- and r_e-range, it is *local* in the aircraft parameters chosen to calculate L_th and also *local* in the assumed W and R. Therefore, $|w^\mathrm{max}| \approx |L_\mathrm{th}^\mathrm{max}|$ cannot be generalized. In addition, the individual noise characteristics need to be considered. As shown in Table 3, *AtlantikSolar* UAV flights in typical weather yield $\sigma_{\hat{w}} \approx 0.4\,\mathrm{m/s}$ and $\sigma_{\hat{L}_\mathrm{th}} \approx 3\,\mathrm{Nm}$. The measurement deviation ΔL_th required for statistical significance is thus much higher than Δw. Also, because $\Delta \boldsymbol{P}_k$ is inversely proportional to $\boldsymbol{R}$ in Eq. (24) the respective state covariance will decrease slower for the same $\boldsymbol{J}_h$. Overall, the analysis indicates that the L_th-measurement is — primarily due to its noise characteristics — of slightly less overall value than $\hat{w}$. As a secondary measurement it can however significantly augment the

thermal estimation process e.g. by resolving the ambiguities of a solely $\hat{w}$-based estimation process.

3.3.5 Simulation Results

To quantitatively assess the benefits of the induced roll moment measurement a *Matlab*-based simulation environment was set up. It is based on previously available open-source work[14] for EKF-based thermal tracking. In our work, it was extended with the measurement of the induced roll moment L_{th} and both an unscented Kalman filter (Wan & Merwe, 2001) and particle filter (Gordon, Salmond, & Smith, 1993). The extended framework is also published as open-source software.[15] The EKF is the standard thermal estimation approach in this paper because it handles the measurement non-linearities in Eqs. (3) and (9) through a computationally efficient 1st-order linearization and thus allows seamless integration on today's low-power autopilots. However, the linearization problems discussed in Section 3.3.4 motivated the implementation of the 2nd-order UKF and the fully non-linear particle filter using standard *Matlab*-libraries. Results for these filters are presented solely to analyze the benefits of considering higher order terms when using $\hat{L}_{\text{th}}$, but the filters were — in the PF case due to high computational demands — not implemented on the aircraft's low-power autopilot.

In the simulation, the typical Gaussian thermal of Figures 12 and 16 with $X^{\text{real}} = [3\,\text{m/s},\ 120\,\text{m},\ 0\,\text{m},\ 0\,\text{m}]$ (location in global coordinates) is used. The simulated *AtlantikSolar* UAV flies at $v_{\text{air}}^{\text{nom}} = 9.6\,\text{m/s}$. The initial state X_0 and more specifically $X_{0,W}$ and $X_{0,R}$ were purposely assigned a noticeable offset (Table 3) from their true value to allow analysis of the filter convergence. To exclude any influence of bad tuning on the results each filter was tuned individually: First, the filter performance was tuned in open-loop mode (i.e. thermal-tracking not active such that the aircraft trajectory and observability for each simulation are the same) over the space $[Q_W^{\min}, Q_W^{\max}] \times [Q_R^{\min}, Q_R^{\max}] \times [Q_r^{\min}, Q_r^{\max}]$. Then, with fixed $\boldsymbol{Q}$, the same simulation was run over $[P_{0,W}^{\min}, P_{0,W}^{\max}] \times [P_{0,R}^{\min}, P_{0,R}^{\max}] \times [P_{0,r}^{\min}, P_{0,r}^{\max}]$. The performance was compared using the accumulated overall residual

$$\zeta = \sum_{t=0}^{t=T} \sum_{i=1}^{i=4} \frac{|X_{t,i} - X_{t,i}^{\text{real}}|}{X_{t,p}^{\text{real}}}\ , \tag{25}$$

where t is the simulation time, i represents the i-th component, and p indicates with which thermal parameter in X^{real} the respective residual is normalized. We chose $p = 2$ for $i \geq 2$ and $p = i$ otherwise. Comparing ζ shows that the filters yield the best performance for approximately the same $\boldsymbol{Q}$, but the UKF's performance is optimized at a lower $\boldsymbol{P}_0$ than for the other filters (Table 3). The UKF is configured with $\alpha = 0.1$, $\beta = 2$, $\kappa = 0$. The PF uses $N = 5000$ particles, and to avoid impoverishement of the particle set the roughening method by Gordon et al. (1993) is implemented with $K = 0.05$.[16] No process- or measurement-noise is applied to the internal simulation states. The measurement covariance $\boldsymbol{R}$ is not zero, but represents a best-case estimate of a calm flight day outside. Overall, the following evaluation is not set up to represent a typical flight day, but to demonstrate the best-case filter performance as *deterministically* as possible.

Figure 17 shows simulation results for the test cases A, C and D of Section 3.3.4. In *Test Case A*, the aircraft approaches from the bottom left and starts tracking as soon as the thermal is detected. The purely w-measurement-based approach yields acceptable tracking performance, but when using both the w- and L_{th}-measurements the EKF centers the thermal faster and the residuals and 1-σ confidence intervals (representing the state covariance) also converge much more quickly. In fact, when only $\hat{w}$ is used both the W- and R-residuals stay constant (with a 35 m radius error) as soon as the thermal is centered because they are not simultaneously observable anymore. The accumulated residuals are $\zeta_{\hat{w}} \approx 813$ and $\zeta_{\hat{w}+\hat{L}_{\text{th}}} \approx 185$. *Test Case C* represents an open-loop (i.e. no active thermal tracking) thermal fly-by. The EKF can directly estimate the relative thermal location using the L_{th}-measurement, but using only the w-measurement it cannot deduce any information in the r_{e}-axis and therefore falsely assumes that the thermal always lies directly ahead or behind the aircraft (which is why the trajectory of the thermal estimate is not visible). The r_{e}-residual

[14] Released by Samuel Tabor at https://github.com/samuelctabor/Soaring_simulation.

[15] Our extensions have been pushed back into the original repository, see the previos footnote.

[16] The filter-specific parameters for both UKF and PF were either obtained from standard literature, or, if in doubt, also optimized within a specific range through running multiple simulations.

and its covariance do not decrease at all, the other residuals decrease similarly fast. An advantage with the L_{th}-measurement is that the aircraft can quickly turn to the correct side. The accumulated residuals are $\zeta_{\hat{w}} = 274$ and $\zeta_{\hat{w}+\hat{L}_{\text{th}}} = 235$. *Test Case D* is also an open-loop test and shows the disadvantages of the EKF's linearization step. We assume that the aircraft approaches from the bottom and that the thermal is (e.g. due to turbulence) initialized on the wrong side of the aircraft. Neither the EKF nor the UKF manage to correct this wrong initialization using $\hat{L}_{\text{th}}$, instead, the EKF even shifts the thermal center further away from the true position because of the Jacobian-properties explained in Figure 16. Employing the 2nd-order UKF also does not help given the high degree of non-linearity.[17] In contrast, the particle filter succeeds to find the true thermal position, radius and core strength. The accumulated residuals are $\zeta_{\text{EKF},\hat{L}_{\text{th}}} \approx 5900$, $\zeta_{\text{UKF},\hat{L}_{\text{th}}} \approx 5500$ and $\zeta_{\text{PF},\hat{L}_{\text{th}}} \approx 900$. The test case demonstrates that the full potential of the induced roll moment measurement can only be exploited with a fully non-linear filtering technique such as the particle

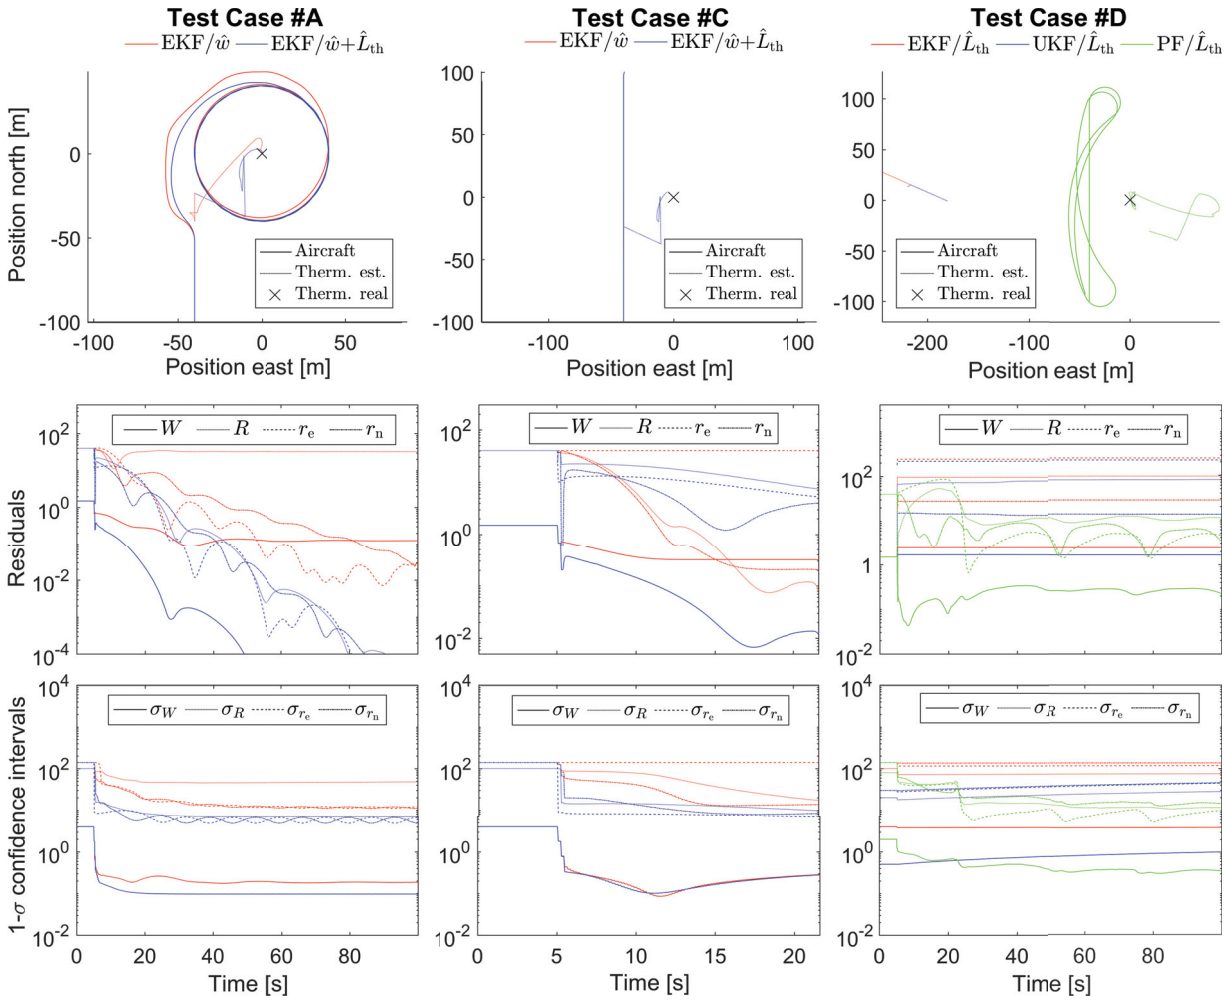

Figure 17: Estimation performance with different measurements ($\hat{w}$ only, $\hat{L}_{\text{th}}$ only, or $\hat{w} + \hat{L}_{\text{th}}$) and filters (EKF/UKF/PF — see top legend) in a simulation of the test cases A, C and D of Section 3.3.4. The top plots show aircraft and estimated thermal position over time. The middle plots show the magnitude of the state residuals, i.e. $|X^{\text{real}} - X|$, and the bottom plots contain the 1-σ (68 %) confidence intervals. Adding the induced roll moment measurement speeds up the filter convergence in all test cases.

[17] To show these issues with the EKF and UKF the w-measurement had to be disabled completely ($R_{\hat{w}} \to \infty$) as it would otherwise correct and eventually allow the filters to find the true thermal position. With the default $R_{\hat{w}}$ of Table 3 such a situation thus does not occur, but the $\hat{L}_{\text{th}}$-based corrections will of course occasionally be counterproductive in an EKF/UKF.

filter.

Figure 18 provides an overall performance comparison of the w- and $w + L_{\text{th}}$-measurements for all three filters. To cover all possible ways of encountering a thermal, the figure is based on $N = 1000$ runs of each filter/measurement-combination in which the aircraft starts with random initial position, heading and pre-programmed trajectory (straight lines, zig-zag, circles). However, because the thermal tracking is active the aircraft soon begins to center the thermal. For the EKF with the w-measurement all residuals decrease at first, but as soon as the thermal is centered sufficiently the W- and R-residuals — as discussed before — do not decrease anymore. With the additional L_{th}-measurement, the residuals decrease faster and are not subject to any unobservability-issues. The state uncertainty of the filter is also lower. The same qualitative result is obtained for the UKF. The particle filter shows a number of particularities: First, in the short period after the thermal detection the residuals converge much quicker than for the EKF and UKF because, due to the high initial uncertainty, the PF spreads its particles over the whole four-dimensional state space and then quickly rejects those that do not fit. However, in the long term the residual convergence slows down. This is explained by the implemented particle roughening which constantly and randomly disturbs the particles and thus also the estimated mean. Due to the relatively large roughening constant K, the particle filter in this paper is designed for thermal estimation robustness rather than precision. Also note

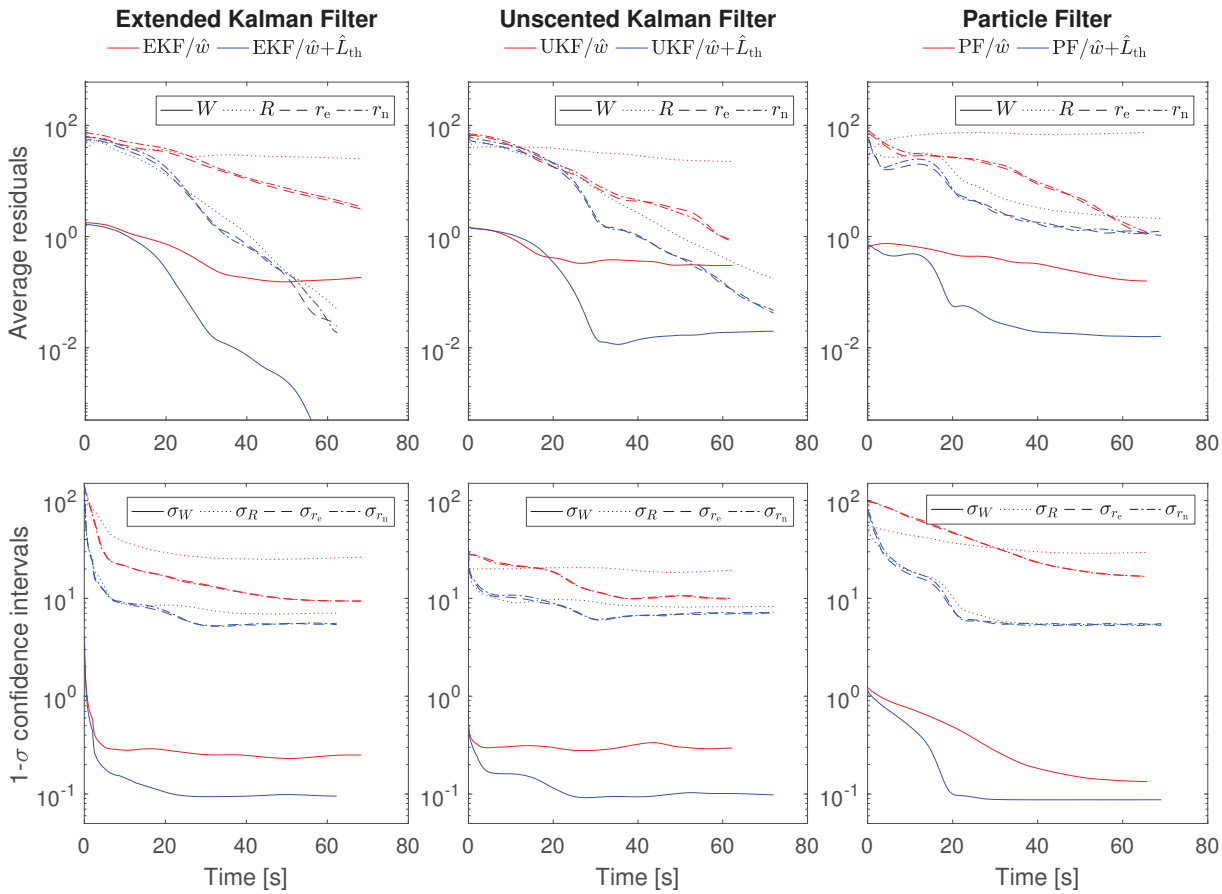

Figure 18: Thermal estimation performance with different measurements ($\hat{w}$ only or $\hat{w} + \hat{L}_{\text{th}}$) and filters (EKF/UKF/PF). The figure shows the residuals $|X^{\text{real}} - X|$ and 1-σ confidence intervals averaged over $N = 1000$ randomized runs of each filter/measurement-combination. All data is plotted from thermal detection ($t = 0\,\text{s}$) until the first simulation run has finished (after 100 s of simulation time). All filters converge faster if the induced roll moment measurement is added.

that when only using $\hat{w}$ the R-estimate diverges[18]. When the L_{th}-measurement is added, the radius and core strength become observable again and the residuals and the uncertainty converge satisfactorily.

Table 4 summarizes the thermal tracking results in the form of the accumulated residuals ζ. When using the *simulation* parameter configuration from Table 3 and only the w-measurement, the UKF performs slightly better than the EKF whereas the diverging R-estimate causes the highest overall residual for the PF. Adding the induced roll moment measurement approximately reduces the accumulated residuals by half. While the EKF and UKF show a similar overall residual[19], the particle filter benefits most from the L_{th}-measurement and clearly shows the lowest accumulated residual overall. For the higher yet more realistic measurement noise of the *flight* parameter configuration the convergence is slower. The accumulated residuals are around 5 % higher if only the w-measurement is used, and because the increase of $R_{\hat{L}_{\text{th}}}$ between the simulation- and flight parameter sets is higher than that of $R_{\hat{w}}$, the overall residuals increase by approximately 10 % for the $\hat{w} + \hat{L}_{\text{th}}$ case. The main conclusion from the simulation results is that *if* the model structure — in this case the Gaussian thermal model — is known well, then the L_{th}-measurement provides very significant potential to speed up the convergence process and to avoid ambiguities in the thermal estimate. Due to the non-linearity in both the w- and L_{th}-measurements fully non-linear estimation techniques such as a particle filter should be used if the computation power is available.

Table 4: Thermal tracking simulation results in the form of the accumulated residuals ζ. The parameter configuration can be retrieved from Table 3.

	Sim. parameters			*Flight parameters*		
Measurements	EKF	UKF	PF	EKF	UKF	PF
$\hat{w}$	1423	1272	1850	1459	1327	1987
$\hat{w}+\hat{L}_{\text{th}}$	760	780	551	808	867	653

4 Flight Results

The goal of the fully-autonomous multi-day flight presented in this section was to test the solar-powered UAV platform and the individual technical contributions of Section 3 under the simulated conditions of a large-scale search-and-rescue mission. To this purpose, the *AtlantikSolar AS-3* UAV was launched autonomously on July 19th 2016 at 16:31 solar time (18:02 local time) at latitude $\Phi = 47°N$ (Hinwil, Switzerland). The mockup mission assumed that one or more humans got lost during an incident. During the whole day and night the aircraft therefore performed low-altitude loitering- and scanning-patterns over the test area (confined to ca. 2 km^2 by flight regulations) to find the humans using its color- and infrared-cameras. Live-imagery was streamed to the ground via WLAN. At the same time, the automatic human detection framework estimated and transmitted the GPS location of the humans to the ground station for forwarding to first-aid relief teams. The two crucial mission challenges were, first, to fly through the night with the continuously operating SAR-payload, and second, to guarantee fully autonomous flight in challenging environmental conditions during day and night. Both objectives were achieved. After crossing the night with sufficient energetic margins, the battery recharging process was supported by enabling autonomous updraft tracking. The aircraft landed autonomously at 18:40 solar time (20:11 local time) on July 20th after 26 hours and 9 minutes of flight and a covered ground distance of 869 km. Not a single pilot stick was moved during the whole flight. The field deployment represents the first time that a low-altitude solar-powered UAV has demonstrated *perpetual endurance* while operating a day/night-capable payload, and is also the first time since suggested

[18] As an explanation, assume that in the particle weighting step the PF needs to match a $\hat{w} = 3\,\text{m/s}$ measurement to all thermal particles. It will give a large weight to strong close thermals independent of their radius (see Figure 12), however, thermals that are further away are only considered if they have a sufficiently large radius R to induce a large w at the aircraft's position. The mean radius R therefore increases. The roughening process contributes by constantly spreading out the particles.

[19] The slightly worse UKF results are unexpected, i.e. literature suggests that the UKF should be superior to the EKF. However, the difference is small and the two filters are actually implemented differently (the EKF was implemented by hand, whereas the UKF uses the *Matlab*-internal implementation), so the difference is considered to lie within the simulation uncertainty.

by Wharington (1998) that an aircraft has combined solar-electrical propulsion and autonomous thermal tracking in flight. A video of the flight is available.[20]

4.1 Energetic Results

Figure 19 visualizes the complete 26-hour flight. The aircraft is launched in clear-sky conditions at 16:31 solar time with fully charged[21] batteries ($SoC = 100\,\%$). At $t = 17.81\,\mathrm{h}$ (all solar times unless stated otherwise), while the aircraft is in low-altitude loitering mode, the solar power generation drops below the required level flight power and the battery discharge begins. Sunset occurs at $t_{\mathrm{ss}} = 19.53\,\mathrm{h}$. The mean power consumption during the night ($19.57\,\mathrm{h} - 28.32\,\mathrm{h}$) is $P_{\mathrm{out}} = 61.3\,\mathrm{W}$, of which the payload consumes $P_{\mathrm{pld}} = 9\,\mathrm{W}$ (and $10\,\mathrm{W}$ if the RGB camera had not been switched off). In addition, from around $t = 26\,\mathrm{h}$ on, steady but severe winds of up to $v_{\mathrm{wind}} = 9\,\mathrm{m/s}$ force the aircraft to fly at $v_{\mathrm{air}} \approx 10.5\,\mathrm{m/s}$ instead of its power-optimal speed $v_{\mathrm{air}}^{\mathrm{opt}} \approx 9.2\,\mathrm{m/s}$. Sunrise occurs at $t_{\mathrm{sr}} = 28.19\,\mathrm{h}$. The minimum state-of-charge of $SoC_{\mathrm{min}} \approx 26\,\%$ is reached at $t \approx 30.23\,\mathrm{h}$ (7.76 h local time) despite carrying and operating the SAR-payload through the complete night. Due to the combined use of solar power and autonomous thermal updraft following, the batteries again reach

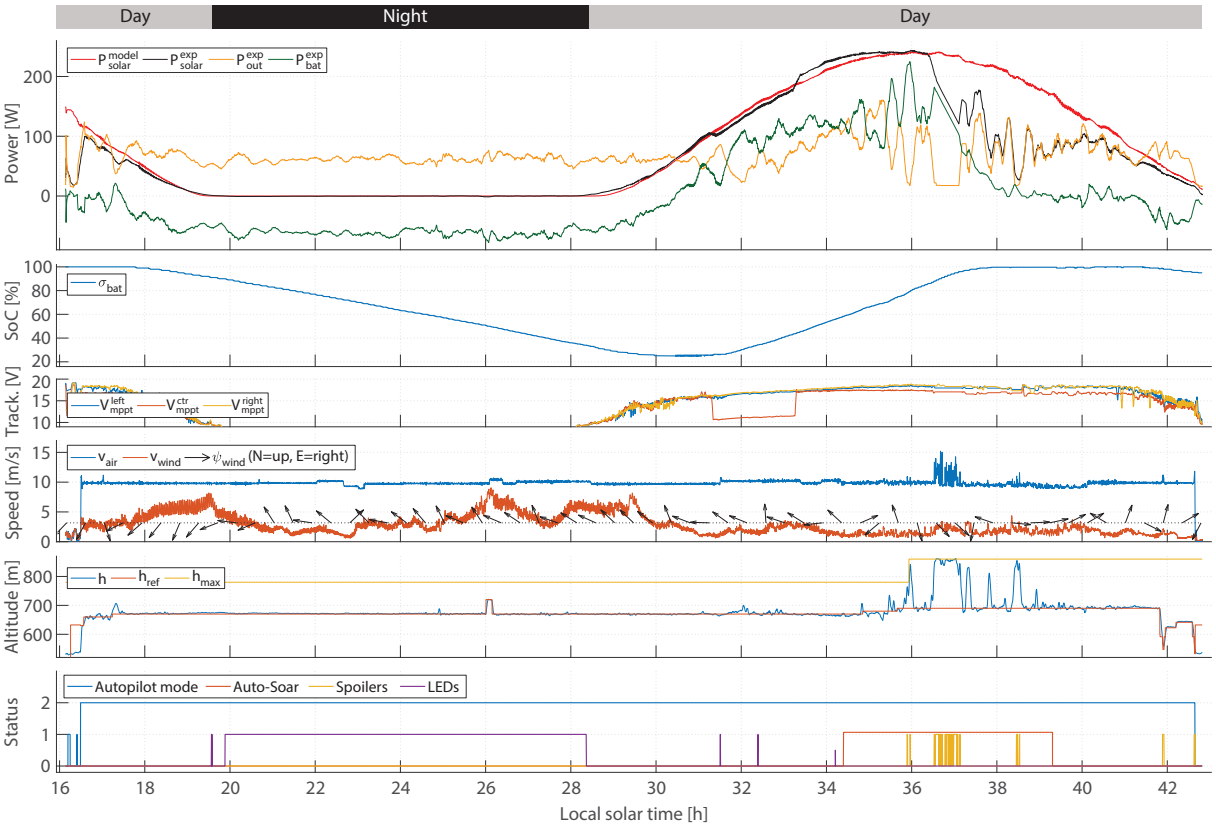

Figure 19: Visualization of the 26-hour continuous solar-powered search-and-rescue flight of *AtlantikSolar* AS-3 in local solar time. The modeled solar power $P_{\mathrm{solar}}^{\mathrm{model}}$ is derived from the *Full Solar Power Model* in (Oettershagen, 2017). The data is sampled at 3 Hz and a two-sided moving average filter with semi window length 200 samples is applied to all data except for the *Status* plot. In the status plot, the *Autopilot mode* flag is [0 = manual flight, 1 = assisted flight, 2 = autonomous flight] and the *Auto-Soar*, *Spoilers* and *LEDs* flags are [0=OFF, 1=ON]. Note that the solar-power income drop at $t = 31.31\,\mathrm{h}$ is due to one of the Maximum Power Point Trackers (MPPT) failing to track the optimum solar panel voltage.

[20] See www.youtube.com/watch?v=8m76Mx9m2nM. Further videos are available at www.youtube.com/user/AtlantikSolar.
[21] The full-battery launch was not planned, but the aircraft had quickly recharged its batteries while waiting for launch.

$SoC = 90\,\%$ at $t_{\mathrm{fc}}^{90\,\%} = 36.64\,\mathrm{h}$, then gradually start limiting the charging power to reduce battery wear, and finally reach $SoC = 100\,\%$ at $t_{\mathrm{fc}} = 37.87\,\mathrm{h}$ solar time. Energetically-perpetual flight is thus achieved. Due to upcoming cloud cover, the discharge on the second day starts at $t = 41.45\,\mathrm{h}$ and thus a bit earlier than on the first day. The charge margins are therefore only a rough estimate, but yield $T_{\mathrm{cm}}^{90\,\%} \approx 4.81\,\mathrm{h}$ and $T_{\mathrm{cm}} \approx 3.58\,\mathrm{h}$. While the autonomous thermal updraft tracking has helped to reduce the recharging time (Section 4.4), the higher aircraft power consumption and larger batteries lead to a significantly increased recharge time with respect to (Oettershagen et al., 2017). The aircraft is landed autonomously at $t = 42.66\,\mathrm{h}$ solar time with $SoC = 95\,\%$ and the hypothetical capability to continue the flight through another night.

Overall, the energetic margins confirm or even slightly exceed those obtained via our conceptual design and analysis framework in Figure 4 ($SoC_{\mathrm{min}} = 22\,\%$ and $T_{\mathrm{cm}} = 3.9\,\mathrm{h}$). The small errors lie within the expected uncertainty due to atmospheric disturbances, the simplifications in our analysis framework and SoC measurement errors. As expected, the energetic margins are significantly lower than those retrieved for the 81-hour flight without payload in our previous work. However, due to the payload development and the platform and flight-autonomy optimizations performed in Section 3, the multi-day search-and-rescue flight with payload could be demonstrated with sufficiently high energetic safety margins.

4.2 Aerial Sensing and Human Detection

Operating over the full 26 hours of the mockup search-and-rescue flight, the aerial sensing system's color- and infrared-camera were used for both aerial mapping and automatic human detection. The aerial mapping results of Figure 20 show a reconstruction from a single circle of a loiter-path via the a) color camera during the day and b) the infrared camera during the night. The mapping process uses the UAV's state estimator output to georeference the captured images and an accurate 3D representation of the terrain is then computed using Pix4D.[22] The 3D map reconstruction is not a novelty per se and not the focus of this paper, but it is provided for completeness and to highlight the benefits of long-endurance day/night aerial sensing. In SAR missions after landslides, earthquakes or fires these maps could easily be used to identify the fire's extent or obstructed roads during both day and night.

The primary focus of this section is to briefly present human detection examples and the algorithm's performance in flight. A sample infrared image and the corresponding visual-spectrum image, showing three humans, is presented in Figure 21. The images were recorded onboard *AtlantikSolar* at 85 m altitude above ground at 11:46 local time. As described in Section 3.1.3, in the first step of the human detection pipeline the blob detector detects potential humans in the infrared image. These detections are visualized as blue circles in Figure 21a. In the second stage of the detection pipeline, the HOG/SVM-IS is used to classify the potential humans. The corresponding threshold and classification results are shown in Figure 22. All detections with a classification output higher than the threshold of 0.5 are classified as humans. The four positive detections correspond to the four red rectangles in Figure 21a. In stage three, the human detections from the infrared spectrum are mapped into the visual spectrum, samples around the detection centers from stage two are generated and the corresponding image patches are classified using HOG/SVM-VS. The final human detections are shown as purple rectangles in Figure 21b. Overall, fusing the visual spectrum with the infrared spectrum results in a considerable improvement. The flight data confirms the findings of our previous publication (Kümmerle et al., 2016), in which a detailed quantitative analysis showed that an improvement of almost 10% in miss rate in the relevant false positive per image (FPPI) can be achieved.

4.3 Autonomous Launch and Landing

4.3.1 Launch Results

Results from the autonomous hand-launch procedure described in Section 3.2.1 and applied during the 26-hour flight are shown in Figure 23. Higher noise can be seen in the x-body acceleration data during the

[22] Pix4D is a commercially available software. See http://www.pix4d.com.

(a) Orthomosaic from color images taken during the day.

(b) Orthomosaic from pseudo-colored thermal images taken during the night.

Figure 20: Orthomosaics created from images recorded during one circle flown with the *AtlantikSolar* UAV.

(a) The blue circles symbolize all blob detections in the IR spectrum. The red rectangles show positive human classifications based on the HOG/SVM-IS.

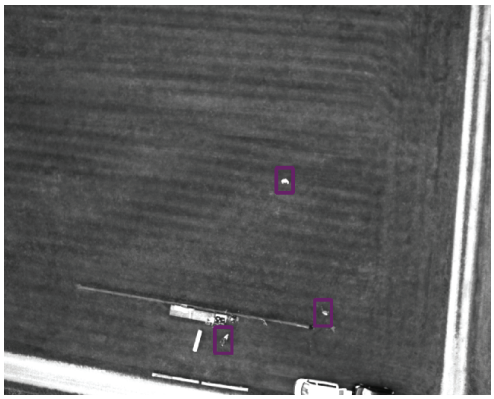

(b) The IR-detections are mapped to the visual spectrum and classified using HOG/SVM-VS. Purple rectangles show the final human detections.

Figure 21: Output of blob and HOG detector applied to an infrared image.

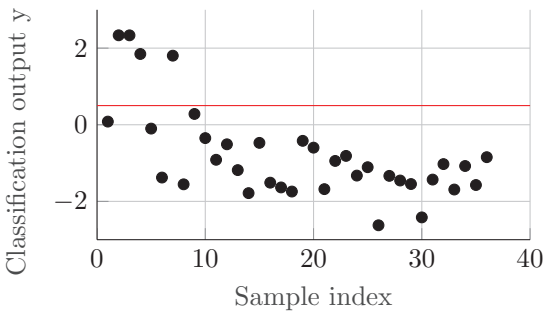

Figure 22: Classification results for the candidates in the infrared spectrum image (Figure 21). The candidates are classified as human if their classification output y is above the threshold (red line).

running start (region R3), pointing to the importance of a sufficiently high acceleration threshold. On the other hand, it can be seen at the beginning of region R4, that the highest spike in x-body acceleration, at launch, may not have passed the acceleration threshold.[23] We observed the threshold of 10 m/s² to be just exceeded in most launches; setting a lower threshold risks false detection during the running start. This motivates the additional launch-detection metric of the airspeed threshold v_{thresh}, which is exceeded even earlier than the actual launch in region R3. At this point, the take-off procedure was initiated in the autopilot, commanding level roll, and using *climbout* TECS logic with the small minimum pitch angle command, $\theta_{\text{min}}^{\text{TO}}$, seen in the pitch plot.

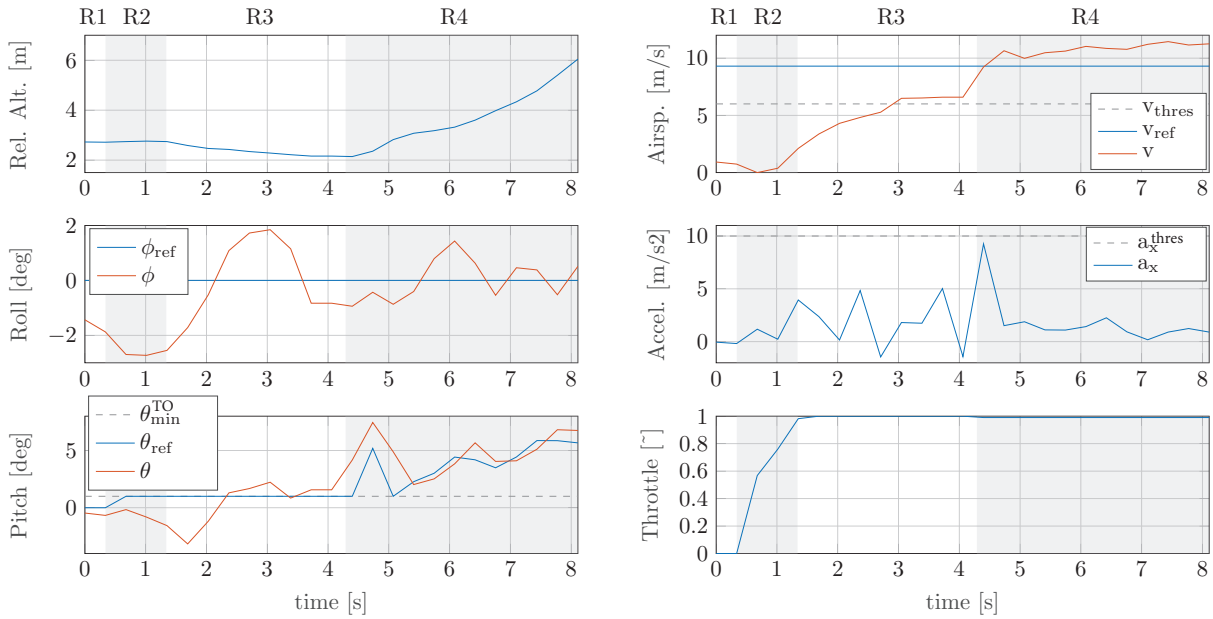

Figure 23: Flight data during hand-launch. Region R1 is prior to engaging the launch-detection module, R2 shows the throttle ramp-in to idling motor ($u_{\text{thr}}^{\text{idle}} = 100\,\%$ for *AtlantikSolar*) and when the hand launcher began his running start, R3 extends from start of running until the launch itself, and R4 is post-launch.

4.3.2 Landing Results

At the end of the flight, the landing procedure described in Section 3.2.2 was implemented, starting with generation of the race-car/spiral-down landing path, and a closely monitored descent within line-of-sight of the ground station until touchdown. The path with entrance/exit waypoints for each segment as well as the UAV's position and terrain estimate for the final approach are shown in Figure 24. Lateral-directional ($\mathcal{L}_1$-based) tracking on straight segments showed, at times, several meters offset during the descent. Though, performance on spiral-down tracking, remained up to typical standards with less than 1 m tracking error. The discrepency in performance between curved vs. straight segments could be due to the nature of the guidance algorithm itself, suited originally for circle tracking in a nonlinear sense (Park et al., 2007). However, straight segment tracking during level-flight did not experience the same steady state errors. Most likely, the degraded high-level performance during descent is due to the roll tracking loops' changed trim characteristics and less reactive tracking due to the deployed spoilers (see for example the sporadic roll tracking offsets in Figure 25). The loop should be retuned to use a different set of gains when the spoilers are engaged. While high-level controllers should typically remain agnostic to low-level functionality, requiring low-level loops to perform well enough to adhere to the hierarchical control paradigm of "sufficient low-level tracking of high-level commands", additional improvements could be incorporated through specific $\mathcal{L}_1$ gains for landing/descent

[23] Note that in all automatic launch and landing results the logging sample rate was 3 Hz due to the limited onboard space for storing logs over the 26 hours of flight. Some higher frequency effects could thus be missed within the data.

or the inclusion of an integral term into the $\mathcal{L}_1$ algorithm.

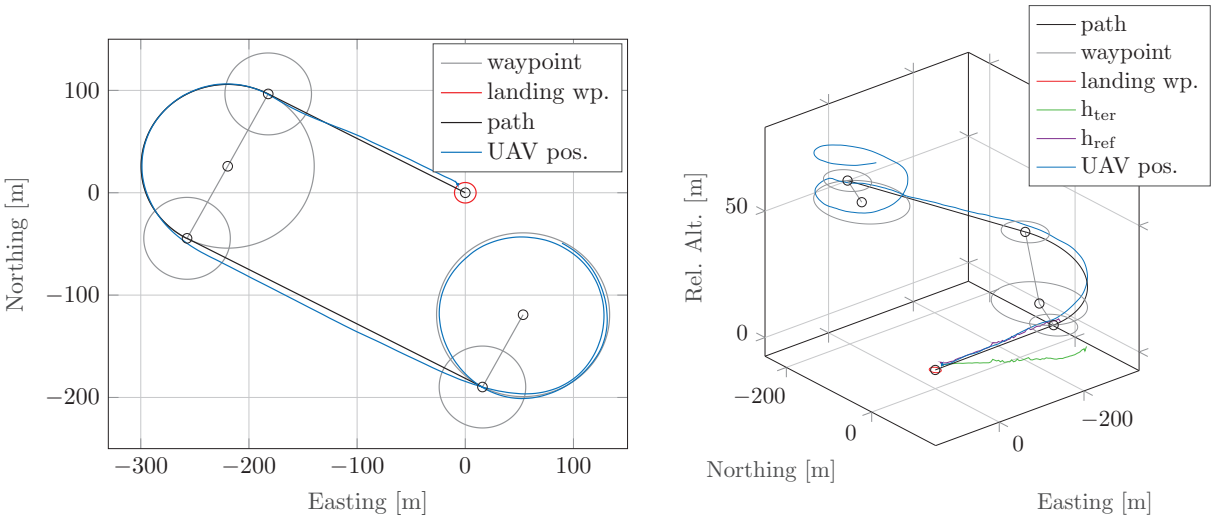

Figure 24: Position tracking during automatic descent and landing (top view - left, 3D view - right). In the right figure, the h_{ter}-curve (green) is the estimated terrain altitude and h_{ref} (purple) is the actual altitude reference generated by shifting the path (black) up by the terrain altitude. The landing waypoint is shown in red with a touchdown radius of 10 m.

Longitudinal tracking of the landing path also saw some offset during parts of the descent, in this case, conversely to the lateral-directional tracking, more so during the spiral-down segments. Again, the likely culprit, despite updated TECS gains, is that the gain-set for the lower level CAS and SAS loops was tuned only for nominal flight but was used for a flight regime outside of its trim condition — i.e.descent with engaged spoilers. TECS also operates on a level-roll assumption, where spiral-down trajectories introduce effects not modeled within the controller. In Figure 25, altitude tracking can be seen to have settled well during the straight segment on final approach. However, the fast change in the landing curve during the flare portion (starting at $h_{\text{fl}} = 1\,\text{m}$, $t = 145.1\,\text{s}$), without more aggressive pitch tuning, resulted in the aircraft continuing along the landing slope for the remaining second of flight before touching down 10 m in front of the landing point. Finally, the importance of the altimeter can also be seen in Figures 24 and 25, where the terrain height estimate h_{ter} measured the ground a couple meters higher than the original home waypoint altitude would have predicted. The discrepancy was likely due to GPS-altitude error, commonly within 4 m, though slight rises or falls in terrain from the point near the ground station where the home waypoint altitude is initialized are also often present. Further, the propeller spin-up seen in the throttle curve at $t = 136\,\text{s}$ is due to the LIDAR's measurement of a slight increase in terrain height (some small bushes) at the edge of the landing strip. Overall, despite minor performance degradation in position tracking during descent, our current configuration allowed precise autolanding within 10 m of the target point. It is expected that integration of the additional functionalities discussed in this section, particularly flight-regime specific low-level tuning, would improve both precision and accuracy further, especially in lateral tracking.

4.4 Thermal Updraft Following

The autonomous thermal updraft tracking described in Section 3.3 was enabled at 10:24 local solar time during the second day of the 26-hour SAR demonstration mission (Figure 19). The aircraft performed constant-altitude loitering while taking advantage of the thermal updrafts that developed over noon and afternoon using the *flight* parameters in Table 3. Thermal tracking was active for 4.9 h until 17:19 local solar time. The algorithm latched into 26 detected updrafts and was able to perform 16 climbs[24] with an

[24] A *subjectively perceived* thermal in Figure 19 can contain multiple climbs (e.g. the large thermal at $t = 36.5\,\text{h}$ contains four climbs of $\Delta h > 50\,\text{m}$ because the thermal broke down at one point but then regained strength).

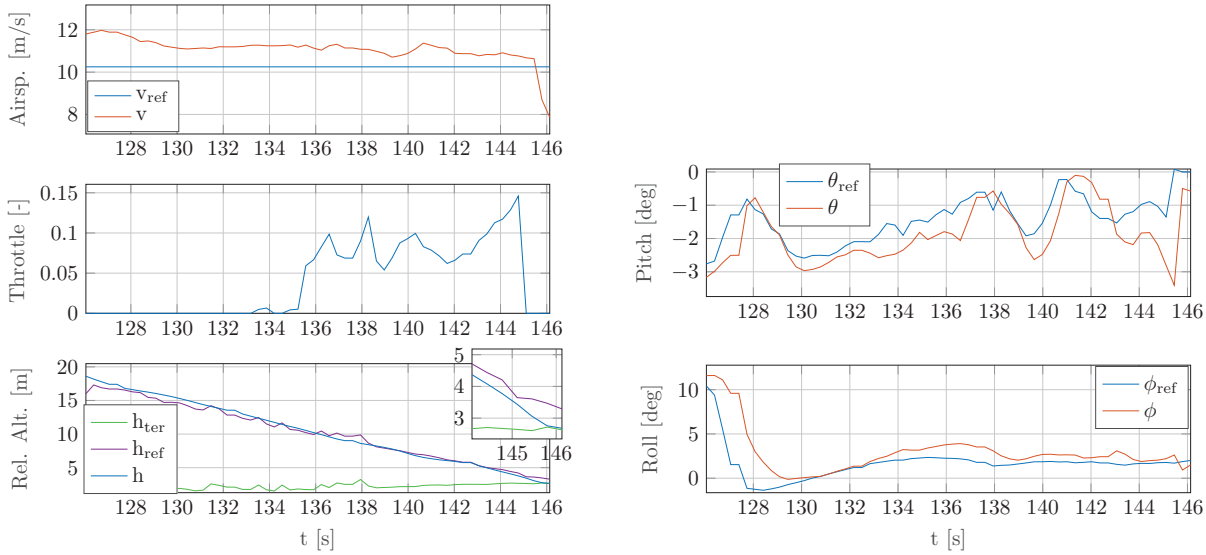

Figure 25: Visualization of the outer loop (left) and inner loop (right) tracking behavior during the final landing approach. Flaring just before touchdown ($t = 145.8s$) is shown in the blown up superposed figure on altitude. The graphs show that despite a solar-powered UAV's challenging flight dynamics and susceptibility to horizontal and vertical gusts, a tracking behavior that allows safe autonomous landings can be achieved.

altitude gain of $\Delta h > 50\,\mathrm{m}$.[25] The average and maximum altitude gain per climb were $\Delta h^{\mathrm{mean}} = 88\,\mathrm{m}$ and $\Delta h^{\max} = 178\,\mathrm{m}$ (in the large thermal at $t = 36.5\,\mathrm{h}$). The UAV climbed in thermals for 1827 s and accumulated 1406 m of unpowered total altitude gain while continuing its loitering mission. Note that the altitude gain per climb was limited by the stringent maximum altitude limit $h_{\max}$ applied because of both camera ground resolution requirements and local airfield restrictions. This altitude limit was reached seven times, in which case the aircraft engaged its spoilers automatically and unlatched from the thermal (see also Figure 26), but consequentially left a significant part of the strong thermals that brought it up to $h_{\max}$ unused. Overall, inside a thermal the average climb rate was $\hat{v}_z = 0.77\,\mathrm{m/s}$ and the average updraft speed was $\hat{w} = 1.20\,\mathrm{m/s}$. The maximum climb speed and updraft speed averaged over one single thermal were $\hat{v}_z = 1.32\,\mathrm{m/s}$ and $\hat{w} = 1.65\,\mathrm{m/s}$ (thermal at $t = 36.5\,\mathrm{h}$). These four values seem low for a mid-summer day at the given latitude. However, the thermals were generally weak ($\hat{v}_z < 1\,\mathrm{m/s}$) in the morning and afternoon, and even over noon the aircraft climb speed measurement is mostly $\hat{v}_z < 2\,\mathrm{m/s}$ and contains no indication that any *strong consistent* updrafts with $W > 2.5\,\mathrm{m/s}$ were present. Nevertheless, overall the thermal tracking allowed 1.45 h of unpowered ($u_{\mathrm{thr}} = 0$) flight — about 32 % of the total time in which thermals were tracked autonomously. Unfortunately, these performance metrics for the *active* thermal tracking cannot be compared in a quantitative way against *passive* thermal compliance flights as for example the 81-hour (Oettershagen et al., 2017) and 28-hour (Oettershagen et al., 2016) flights. In general, the thermal updraft situation varies too significantly from day to day to draw reliable performance conclusions, and in addition both the aircraft configuration and environment for the aforementioned flights were different.

To allow an in-depth thermal tracking analysis, Figure 26 shows the aircraft- and thermal-trajectories in a single thermal climb and Figure 27 visualizes the corresponding measurements, states and state covariances. Note that this specific thermal was not part of the 26-hour flight but from a later flight at a different airfield: As this airfield allowed for a larger altitude band ($h_{\max} - h_{\min}$), the filter had more time to converge and thus a better analysis of the tracking algorithm's convergence can be performed. During $t = [0, 106]\,\mathrm{s}$, the aircraft performs a standard constant-altitude loitering mission in `SEARCH/CRUISE` mode. Both $\hat{v}_z$ and $\hat{w}^{\mathrm{f}}$ indicate downdrafts and, through their fluctuations, also the high level of atmospheric turbulence that makes the multi-step low-pass filtering of the updraft speed (unfiltered $\hat{w}$, mildly filtered $\hat{w}^{\mathrm{f}}$, and heavily filtered

[25] Given the aerodynamic efficiency of the *AtlantikSolar* UAV this represents a no-motor glide time of ca. 2 minutes.

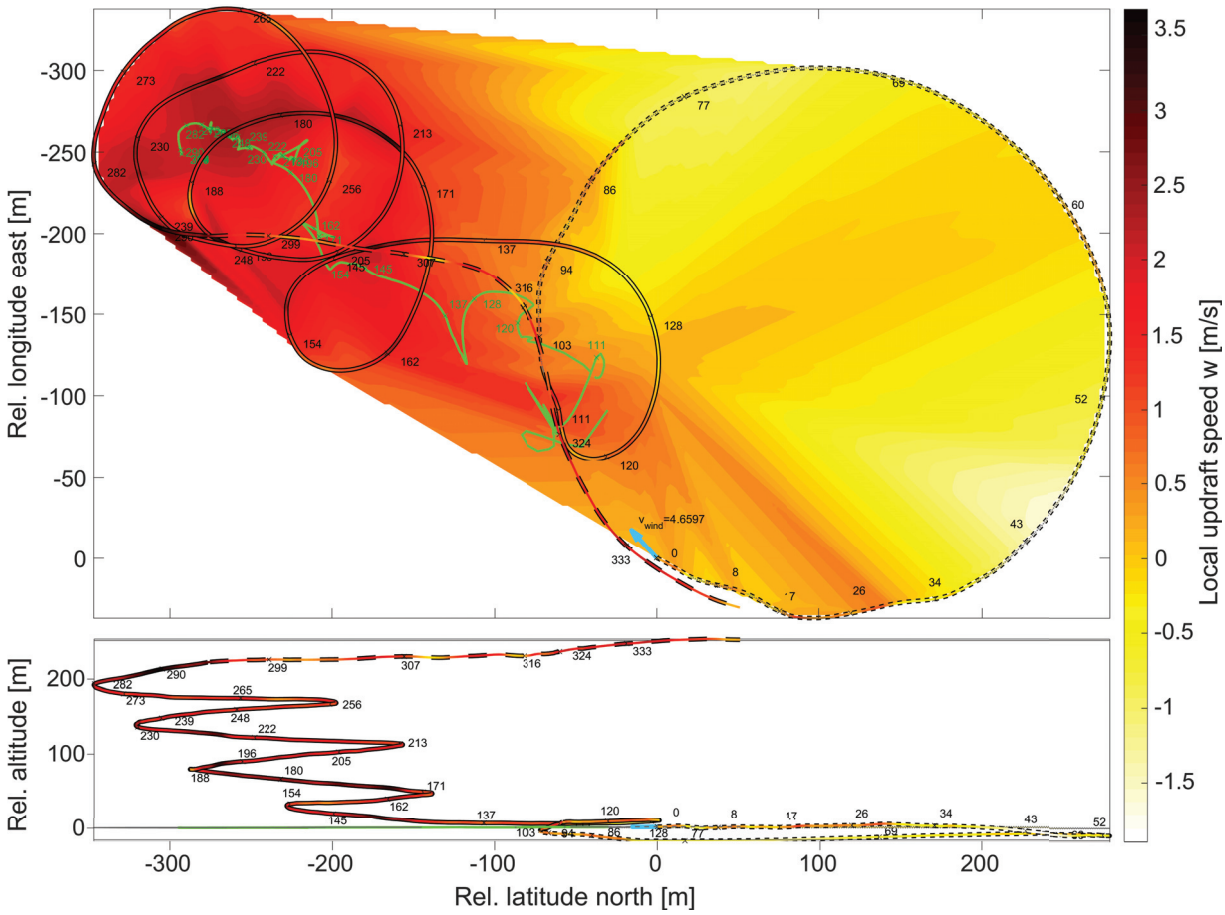

Figure 26: Top-view (top) and side-view (bottom) of the aircraft trajectory (black dotted-, solid-, dashed-lines represent `SEARCH/CRUISE`-, `PRESOAR/SOAR`-, `AVOIDSOAR`-mode respectively) and estimated thermal center trajectory (green) in autonomous updraft tracking. The labels indicate the time in seconds. The aircraft trajectory color represents the local updraft velocity $\hat{w}^{\mathrm{f}}$. The colored surface is retrieved from spatial interpolation of $\hat{w}^{\mathrm{f}}$. The blue arrow represents wind speed and direction.

$\hat{w}^{\mathrm{ff}}$) necessary. At $t = 90\,\mathrm{s}$ the aircraft approaches an updraft, and the increase of $\hat{w}^{\mathrm{ff}}$ triggers the `PRESOAR` mode at $t = 106\,\mathrm{s}$. The Kalman filter initially estimates the thermal center north of the aircraft. After verifying the latching conditions it activates the `SOAR` mode at $t = 113\,\mathrm{s}$ and engages into a left-turn. The turn however leads the aircraft into a region of lower $\hat{w}^{\mathrm{f}}$, and the EKF therefore correctly shifts the estimated thermal center south. Note that despite the first part of the unlatch condition ($\hat{w}^{\mathrm{ff}} < {}^{1}\!/_{4} w_{\mathrm{latch}}$ for $5 \cdot T_{\mathrm{latch}}$) being fulfilled at $t = 130\,\mathrm{s}$, the tracking algorithm does not unlatch, i.e. it provides robustness against wrong initial estimates through the second condition ($T_{\mathrm{latched}} > T_{\mathrm{latch}}^{\mathrm{min}}$) shown in Figure 15. The interplay of the autopilot following the estimated thermal center towards the south and the increasing $\hat{w}^{\mathrm{f}}$ there leads to a quick convergence of the thermal position. At $t = 175\,\mathrm{s}$, the algorithm has lead the aircraft into the vicinity of the thermal core where $\hat{v}_z \approx 1.3\,\mathrm{m/s}$ at $\hat{w}^{\mathrm{f}} \approx 1.6\,\mathrm{m/s}$. The EKF's prediction error or *innovation* (see Eq. (19)) is low and the state uncertainty (represented as 1-σ confidence bounds in Figure 27) has decreased significantly. As a result, the EKF now also predicts the induced roll moment well: Over $t = [175\,\mathrm{s}, 293\,\mathrm{s}]$ we on average predict $L_{\mathrm{th}} = 0.77\,\mathrm{Nm}$ and measure $\hat{L}_{\mathrm{th}}^{\mathrm{f}} = 0.61\,\mathrm{Nm}$. A significant amount of noise that motivates the choice of $\sigma_{\hat{L}_{\mathrm{th}}} = 3\,\mathrm{Nm}$ in Table 3 is however present due to turbulence and the neglected roll acceleration in Eq. (15). In the context of Eq. (2) it is also visible that the horizontal thermal shift over time and altitude aligns with the wind vector v_{wind}. However, the available data does not permit a final conclusion on whether this is due to wind shift or just a result of the estimation process. Finally, after a climb of 189 s the aircraft

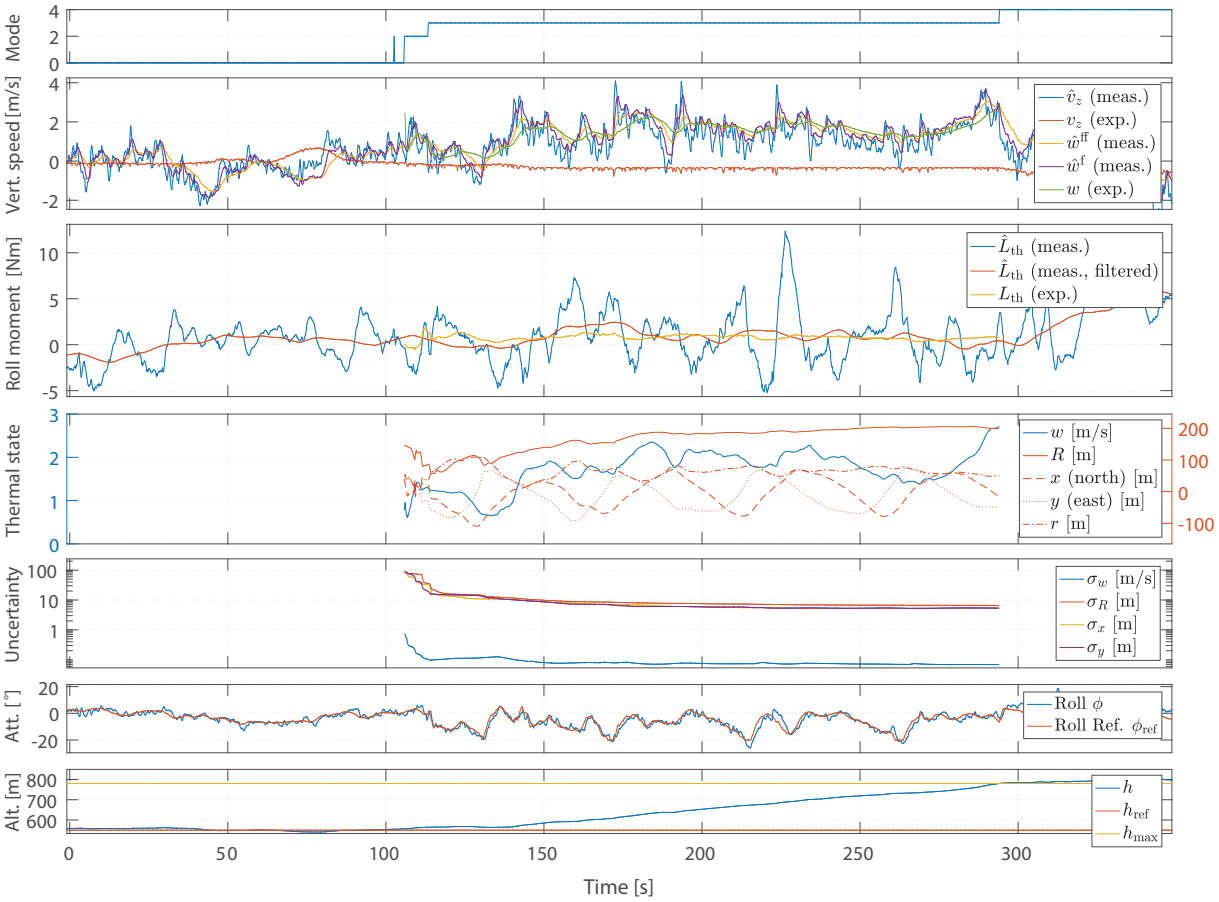

Figure 27: Current updraft tracking mode (0=`CRUISE`, 1=`SEARCH`, 2=`PRESOAR`, 3=`SOAR`, 4=`AVOIDSOAR`), measurements, estimated thermal state, uncertainty in the form of the 1-σ (68%) confidence intervals and aircraft states during the autonomous tracking of the thermal updraft in Figure 26. The data was recorded at 3 Hz. Note that $\hat{L}^{\mathrm{f}}_{\mathrm{th}}$ is the roll moment measurement used on-board the aircraft, and $\hat{L}^{\mathrm{ff}}_{\mathrm{th}}$ is further low-pass filtered a-posteriori (using a two-sided moving average filter with semi window length 100 samples) to allow direct comparison to L_{th}.

reaches the user-specified altitude limit h_{max}. It consequentially deploys its spoilers and unlatches from the thermal by engaging the `AVOIDSOAR` mode. After gaining $\Delta h = 259\,\mathrm{m}$ (enough for a 13 minute unpowered glide) using its autonomous thermal updraft tracking functionality, the aircraft automatically returns to its loitering circle and continues its aerial sensing mission.

The quantitative assessment of the filter convergence improvements due to the induced roll moment L_{th} cannot be performed using the state residuals (as in Section 3.3.5) because no ground truth data exists. Instead, the effect is assessed more qualitatively by comparing measurements and measurement residuals from flights inside and outside of thermals: Figure 28 plots the required aileron command u_{ail} to keep the reference roll angle ϕ_{ref} inside a turn (left) and the expected and measured induced roll moments L_{th} and $\hat{L}_{\mathrm{th}}$ (right). To take the strong noise and different environmental conditions into account, the data is averaged over seven *AtlantikSolar* flights that contain nine sufficiently large and distinct thermal updrafts (six to the left and three to the right of the aircraft). For every thermal, we first wait for filter convergence and thermal-centering such that L_{th} can be estimated from the state X reliably, and then only the remaining interval until $h \geq h_{\mathrm{max}}$ is used (e.g. $t = [175\,\mathrm{s}, 293\,\mathrm{s}]$ in Figures 26 and 27). The 969 seconds of data in which the filter is converged comprise $\Delta h = 1053\,\mathrm{m}$ altitude gain. The comparison data stems from 20 h

of flight in calmer air without thermal updrafts. The required aileron command u_{ail} in Figure 28 (left) confirms Section 3.3.2, i.e. the low dihedral does indeed make *AtlantikSolar* unstable around the roll axis and requires to apply the aileron trim equation (Eq. (16), visualized by the black sine wave) to counteract the tendency to *lean into* the turn — both inside and outside of thermals. Second, the aircraft clearly requires less counteractive aileron while loitering around a thermal, an effect that is caused by the thermal-induced roll moment. Averaged over both left- and right-turns, 1.5 % less counteractive aileron are required. As expected, this is less than the maximum of 6 % calculated for the typical thermal in Section 3.3.2. While subject to significant noise due to atmospheric turbulence, the L_{th}-plots in Figure 28 (right) show that the aileron-trim works: When not in a thermal the average $\hat{L}_{\text{th}}$ is very close to zero, i.e. it is only $\hat{L}_{\text{th}}^{\text{left}} = 0.04\,\text{Nm}$ and $\hat{L}_{\text{th}}^{\text{right}} = -0.01\,\text{Nm}$ when averaged separately for left- and right-turns. Inside thermals, the induced roll moment is clearly visible again. The respective measurements are $\hat{L}_{\text{th}}^{\text{left}} = 0.47\,\text{Nm}$ and $\hat{L}_{\text{th}}^{\text{right}} = -0.65\,\text{Nm}$. Averaged over both left- and right-turns the thermals thus cause an induced roll moment of $|\hat{L}_{\text{th}}| = 0.56\,\text{Nm}$ that can be exploited to improve the thermal estimation process. However, this $\hat{L}_{\text{th}}$ amounts to only 40 % of the $|L_{\text{th}}| = 1.41\,\text{Nm}$ expected by our framework.[26]

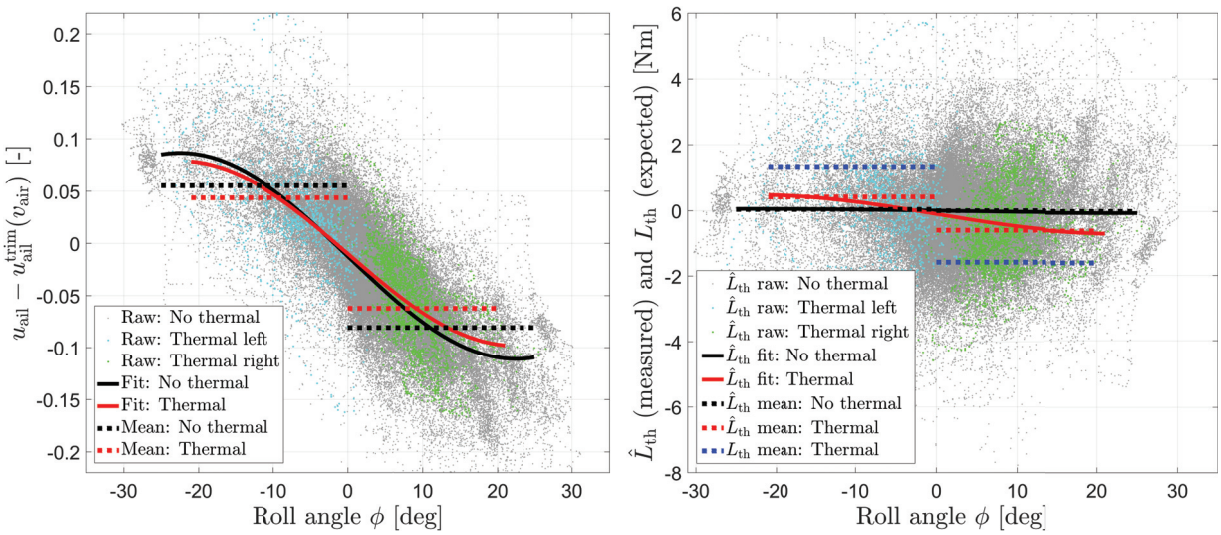

Figure 28: Roll behavior and induced roll moment in *AtlantikSolar* flight testing: Aileron command u_{ail} (where $u_{\text{ail}} > 0$ indicates a roll-right command) vs. roll angle ϕ (left) and induced roll moment L_{th} vs. roll angle (right). The raw data is recorded while the aircraft tracks different roll angle references ϕ_{ref} outside of thermals (grey) and with an estimated thermal on its left (cyan) and right (green) side. The solid lines represent a fit to the raw data according to the trimming relationship in Eq. (16). The dotted lines are piece-wise averages that are calculated separately for thermals to the left and right side of the aircraft. The plots clearly indicate that the roll moment induced by a thermal exists, but is weaker than expected.

Figure 29 provides further details by plotting the measurements and states as a function of the distance r to the estimated thermal core. The induced roll moment is again clearly visible but is subject to significant deviations from its mean. Averaged over r we retrieve $\hat{L}_{\text{th}}^{\text{left}} = 1.10\,\text{Nm}$ and $\hat{L}_{\text{th}}^{\text{right}} = -1.05\,\text{Nm}$ for thermals on the aircraft's left and right side respectively. This represents 78 % of the expected induced roll moments $L_{\text{th}}^{\text{left}} = 1.36\,\text{Nm}$ and $L_{\text{th}}^{\text{right}} = -1.40\,\text{Nm}$ and confirms that the induced roll moment is weaker than predicted. A first explanation may be that the assumed Gaussian updraft speed distribution is a significant simplification. The gradient-based induced roll moment measurement is of course very sensitive to errors in this distribution, which is a general drawback of the proposed method. However, in this particular case Figure 29 (top) shows that the average $\hat{w}$ over the nine thermals does fit a Gaussian distribution with $W = 2.3\,\text{m/s}$ and $R = 100\,\text{m}$ very closely. Another potential reason are errors in the prediction- and measurement-parameters

[26] This measurement residual of course only persists because the noise in $\hat{L}_{\text{th}}$ required to chose a high $R_{\hat{L}_{\text{th}}}$ and the results thus resemble an *open-loop* filtering approach with respect to L_{th}. However, for perfect aircraft-parameter and thermal-model knowledge we still expect that the w-measurement causes $\lim_{t\to\infty} \Delta L_{\text{th}} = 0$. This is clearly not the case here.

of Eqs. (9) and (15). We expect only minor errors in $C_{L_{\text{ail}}}$ and C_{L_p} because they stem from system identification processes. However, the methods proposed to calculate $\frac{dc_\text{L}}{d\alpha}$ neglect a) the non-linearities of the lift-curve at the high α_{nom} of our solar-powered UAV and b) calculate only the mean $\frac{dc_\text{L}}{d\alpha}$ over the wing and thus neglect e.g. that the additional lift Δc_L due to Δw may at the wing tips be reduced due to wing tip vortices. Third, the data from the nine analyzed thermal updrafts may — given the high- and low-frequency noise in $\hat{L}_{\text{th}}$ — not even be enough to allow a final conclusion on the exact magnitude of the induced roll moment. More specifically, the small admissible h_{max} limits the time t available for filter convergence and results in state estimate X errors that can cause equivalent errors in L_{th}. The flights were also performed at low altitude for which thermals are not fully developed yet (Allen, 2006). To remedy these issues, longer flights at higher altitudes are required. The corresponding w-measurements can then be used to verify and potentially fit more elaborate updraft speed distribution models (Allen, 2006) and the L_{th}-measurements can be used to correct $\frac{dc_\text{L}}{d\alpha}$.

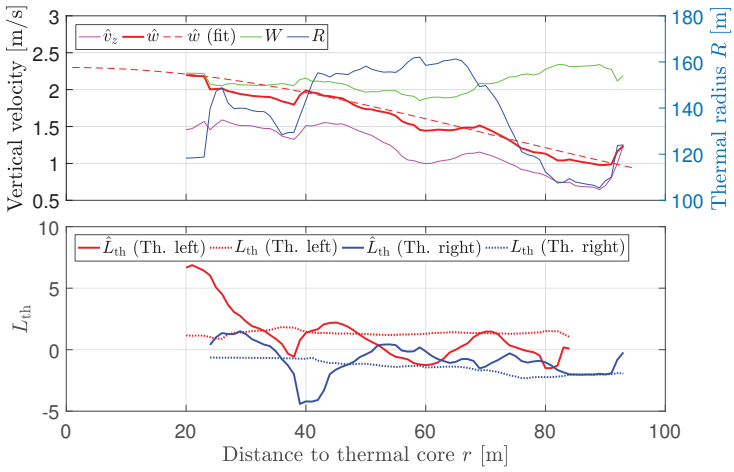

Figure 29: Estimated thermal states (W,R), measurements ($\hat{v}_z$,$\hat{w}$,$\hat{L}_{\text{th}}$) and predictions (L_{th}) vs. the current estimated distance r to the thermal core. The data is averaged over nine different thermal updrafts. The lower figure shows the induced roll moment separately for left- and right-turns around thermals. The dashed red line represents a fit to $\hat{w}$ using Eq. (3) with $W = 2.3\,\text{m/s}$ and $R = 100\,\text{m}$. The plots show that the updraft speed measurements match the assumed updraft speed distribution closely. The thermal-induced roll moment is again visible, however, it is subject to significant measurement noise.

To conclude, on one hand, the induced roll moment measurement leads to much faster filter convergence and is thus highly advantageous *if* good aircraft- and thermal-models are available. The induced roll moment has been demonstrated to exist in real flights and can be exploited for updraft tracking. Our thermal estimation framework allows the end-to-end integration of this L_{th}-measurement in flight. On the other hand, the approach assumes a high-fidelity aircraft model and is thus mostly suited for advanced users. Due to its gradient-based nature, the L_{th}-measurement is sensitive to errors in the updraft distribution model and to turbulence — which together result in noticeable noise. The flight tests show that the *magnitude* of the L_{th}-measurement is lower than predicted. Further flight campaigns to better identify uncertain aircraft parameters and the thermal updraft distribution are suggested. Because of these factors, the induced roll moment measurement should mainly be employed as a *secondary* measurement that extends the *primary* updraft speed measurement by resolving its demonstrated ambiguities.

5 Conclusion

This paper has presented and assessed the technological contributions that have enabled the first-ever fully-autonomous *perpetual endurance* flight of a small solar-powered UAV with a day/night-capable sensing

payload. In the context of a 26-hour long search-and-rescue mockup mission it is demonstrated that, first, the developed sensing system and algorithms can detect humans automatically and reliably during day and night despite stringent weight- and power-requirements from the UAV platform. Second, the solar-powered UAV platform optimizations have allowed to achieve this day/night-flight with payload with sufficient energetic margins. Third, despite the challenging flight characteristics of solar-powered UAVs, this paper has extended existing functionality for autonomous launch and landing and applied it to a *perpetual flight*-capable solar-powered UAV for the first time. Fourth, an autonomous thermal updraft tracking framework that allows the end-to-end integration of the thermal-induced roll moment has been formulated and implemented on the aircraft's low-power autopilot. The main *Lessons Learned* in each of these areas are:

- *Payload development and platform optimization*: Despite recent technological advances in batteries and solar cells, energetic considerations are still dominant if perpetual flight with payload shall be achieved. Day/night sensing operations as demonstrated in this paper still always require, first, the use of lightweight but even more importantly power-saving payloads and, second, the explicit optimization of the UAVs battery system to that payload.
- *Automatic launch and landing*: The highly optimized aerodynamic design of solar-powered UAVs inherently comes with a susceptibility to environmental factors such as gusts or thermal updrafts. These disturbances cannot easily be predicted and, as seen multiple times during the extensive flight testing for this paper, often also cannot be corrected in time due to the slow dynamics of solar-powered UAVs. Every automatic launch and landing framework for solar-powered UAVs thus requires comprehensive safety checks and contingency measures (e.g. go-around functionality during landing if safety-tolerances are exceeded) as detailed in this paper.
- *Autonomous thermal updraft tracking framework*: Due to their efficient aerodynamics, solar-powered UAVs can greatly benefit from autonomous updraft tracking algorithms. Their usually large wing span also allows them to benefit more than other UAVs from fusing the thermal-induced roll moment measurement $\hat{L}_{\mathrm{th}}$. The simulations performed for this paper clearly demonstrate the benefits of using $\hat{L}_{\mathrm{th}}$ and flight tests prove that the effect exists in reality. However, the effect is weaker than expected and subject to significant noise. The *Lesson Learned* is that the updraft speed should remain the primary measurement, but it should be augmented with the induced roll moment as a secondary measurement.

All combined, the contributions in this paper have allowed a 26-hour aerial sensing flight in which the safety pilots did not need to touch the control sticks once. The individual technical contributions of this paper are considered core functionality to guarantee ease-of-use, effectivity and reliability in future multi-day aerial sensing operations with small solar-powered UAVs.

Acknowledgments

This work was supported by Armasuisse contract #043-12, the EU-FP7 research project SHERPA and a number of project partners and generous individuals. We'd like to thank Rainer Lotz and Adrian Eggenberger for acting as safety pilots, Dr. Stefan Leutenegger, Dr. Igor Gilitschenski and Dr. Juan Nieto for their feedback on the paper, Julius Kümmerle for help in developing the human tracking framework, and all project partners[27] for their significant contributions to the project.

[27] A list of project partners is available at http://www.atlantiksolar.ethz.ch/?page_id=187

Appendices

A Derivation of Load-factor Correction

We designate the measured aircraft sink-curve or descent-rate in level flight with 1 g vertical acceleration as $v_\mathrm{d}(v_\mathrm{air})$[28]. Let's assume our *current* flight condition is $n_0 = 1$ at the fixed airspeed $v_{\mathrm{air},0}$ and descent rate $v_{\mathrm{d},0} = v_\mathrm{d}(v_{\mathrm{air},0})$. Then, the lift force $F_\mathrm{L} = mg$ where m is the total aircraft mass and g is the gravitational constant. Combining the standard lift- and drag-force equilibrium equations and $F_\mathrm{D} = P_\mathrm{D}/v_\mathrm{air} = mgv_\mathrm{d}/v_\mathrm{air}$ (where P_D is the power loss to drag) we can formulate

$$c_\mathrm{L}(v_\mathrm{air}) = \frac{2F_\mathrm{L}}{\rho v_\mathrm{air}^2 S} = \frac{2mg}{\rho v_\mathrm{air}^2 S} \tag{26}$$

$$c_\mathrm{D}(v_\mathrm{air}) = \frac{2F_\mathrm{D}}{\rho v_\mathrm{air}^2 S} = \frac{2mgv_\mathrm{d}(v_\mathrm{air})}{\rho v_\mathrm{air}^3 S}\ , \tag{27}$$

where ρ is the air density and S is the wing area. In other words, we have converted the sink curve $v_\mathrm{d}(v_\mathrm{air})$ which is specific to the $n = 1$ level-flight condition to a $c_\mathrm{D}(c_\mathrm{L})$ relationship that is an aerodynamic characteristic of the wing and valid for every aircraft mass or load factor. To retrieve $c_{\mathrm{D},0}$ and $c_{\mathrm{L},0}$ for the current airspeed we plug $v_\mathrm{air} = v_{\mathrm{air},0}$ and $v_\mathrm{d}(v_\mathrm{air}) = v_{\mathrm{d},0}$ into the right-hand side (RHS) of Eqs. (26) and (27). Now assume the aircraft enters a coordinated turn. Our framework assumes that the airspeed is held constant ($v_{\mathrm{air},1} = v_{\mathrm{air},0}$), but then the load factor must increase to n_1 such that $c_{\mathrm{L},1} = n_1 \cdot c_{\mathrm{L},0}$. We plug $c_{\mathrm{L},1}$ into the left-hand side (LHS) of Eq. (26) and solve for the *virtual* airspeed — i.e. $v'_{\mathrm{air},0} = v_{\mathrm{air},0}/\sqrt{n_1}$ — for which this $c_{\mathrm{L},1}$ occurs. Inserting this virtual airspeed into the RHS of Eq. (27) yields

$$c_{\mathrm{D},1} = c_\mathrm{D}\left(\frac{v_{\mathrm{air},0}}{\sqrt{n_1}}\right) = \frac{2mgv_\mathrm{d}\left(\frac{v_{\mathrm{air},0}}{\sqrt{n_1}}\right)}{\rho\left(\frac{v_{\mathrm{air},0}}{\sqrt{n_1}}\right)^3 S}\ . \tag{28}$$

We by now know both the lift- and drag-coefficient for the changed load factor. Finally, we can retrieve the new descent rate $v_{\mathrm{d},1}$ by plugging $c_{\mathrm{D},1}$ (i.e. the RHS of Eq. (28)) into the LHS of Eq. (27) — but now using the actual airspeed $v_{\mathrm{air},1} = v_{\mathrm{air},0}$ instead of the virtual airspeed $v'_{\mathrm{air},1}$ — and solving for $v_{\mathrm{d},1}$:

$$v_{\mathrm{d},1} = \frac{\rho v_{\mathrm{air},1}^3 c_{\mathrm{D},1} S}{2mg} = n_1^{3/2} v_\mathrm{d}\left(\frac{v_{\mathrm{air},0}}{\sqrt{n_1}}\right) \tag{29}$$

B Derivation of EKF Measurement Jacobian

We employ the already introduced $r = \sqrt{r_\mathrm{n}^2 + r_\mathrm{e}^2}$ and in addition define

$$\zeta = \exp\left(-\frac{r_\mathrm{n}^2 + r_\mathrm{e}^2}{R^2}\right) \tag{30}$$

$$\Omega = \frac{1}{24}\frac{dc_\mathrm{L}}{d\alpha}\rho v_\mathrm{air} c_\mathrm{wing} B^3 \cos(\phi) \tag{31}$$

for simplicity. Using the definition of the measurement Jacobian $\boldsymbol{J}_h = dh/dX$ and the definitions of the measurement function $h(X)$ in Eq. (23) and the thermal state X in Eq. (1) we retrieve

$$\boldsymbol{J}_h = \frac{2\zeta}{R^2}\begin{bmatrix} \frac{1}{2}R^2 & -\Omega\left[\cos(\psi)r_\mathrm{e} - \sin(\psi)r_\mathrm{n}\right] \\ \frac{Wr^2}{R} & 2\Omega W\frac{R^2-r^2}{R^3}\left[\cos(\psi)r_\mathrm{e} - \sin(\psi)r_\mathrm{n}\right] \\ -Wr_\mathrm{n} & \Omega W\left[\sin(\psi) + \frac{2r_\mathrm{n}}{R^2}(\cos(\psi)r_\mathrm{e} - \sin(\psi)r_\mathrm{n})\right] \\ -Wr_\mathrm{e} & \Omega W\left[-\cos(\psi) + \frac{2r_\mathrm{e}}{R^2}(\cos(\psi)r_\mathrm{e} - \sin(\psi)r_\mathrm{n})\right] \end{bmatrix}^T\ . \tag{32}$$

[28] For the following derivation, the index *0* represents a specific value not yet corrected for load factor, the index *1* represents a specific value that is corrected for load factor, and without index a variable represents a general functional relationship.

References

Ackerman, E. (2013). Giant Solar-Powered UAVs are Atmospheric Satellites. IEEE Spectrum. http://spectrum.ieee.org/automaton/robotics/aerial-robots/giant-solar-powered-uavs-are-atmospheric-satellites.

AeroVironment. (2013). AeroVironment Solar-Powered Puma AE Small Unmanned Aircraft Achieves Continuous Flight for More Than Nine Hours. Press release. http://www.avinc.com/resources/press_release/aerovironment-solar-powered-puma-ae-small-unmanned-aircraft-achieves-contin.

Allen, M. J. (2006). Updraft Model for Development of Autonomous Soaring Uninhabited Air Vehicles. In *44th AIAA Aerospace Sciences Meeting and Exhibit*.

Allen, M. J. (2007). *Guidance and Control of an Autonomous Soaring UAV*. NASA Dryden Flight Research Center.

Bencatel, R. (2010). Thermal localization. In *International Conference on Autonomous and Intelligent Systems*.

Bower, G. C., Flanzer, T. C., Naiman, A. D., & Saripalli, S. (2010). Dynamic Environment Mapping for Autonomous Thermal Soaring. In *AIAA Guidance, Navigation and Control Conference*.

Bradski, G. et al. (2000). The OpenCV Library. *Doctor Dobbs Journal*, *25*(11), 120–126.

ByeAerospace. (2015). Industry First: Solar-Electric Silent Falcon Prepares for Initial Customer Orders. Press release. Retrieved from http://www.byeaerospace.com/news/2016/3/2/industry-first-solar-electric-silent-falcon-prepares-for-initial-customer-orders

Cheng, M.-M., Zhang, Z., Lin, W.-Y., & Torr, P. H. S. (2014). BING: Binarized Normed Gradients for Objectness Estimation at 300fps. In *CVPR* (pp. 3286–3293). IEEE Computer Society.

Cocconi, A. (2005). *AC Propulsion's solar electric powered SoLong UAV*. AC Propulsion. Retrieved from https://archive.org/details/ACPropulsionSolongUAV2005.

Colella, N. & Wenneker, G. (1996). Pathfinder: Developing a solar rechargeable aircraft. *IEEE Potentials*, *15*(1), 18–23.

Dalal, N. & Triggs, B. (2005). Histograms of oriented gradients for human detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. San Diego, USA.

Davis, J. W. & Keck, M. A. (2005). A two-stage template approach to person detection in thermal imagery. *WACV/MOTION*, *5*, 364–369.

Davis, J. W. & Sharma, V. (2007). Background-subtraction using contour-based fusion of thermal and visible imagery. *Computer Vision and Image Understanding*, *106*(2), 162–182.

Edwards, D. J. (2015). *Autonomous Locator Of Thermals (ALOFT) Autonomous Soaring Algorithm* (tech. rep. No. NRL/FR/5712–15-10,272). Naval Research Laboratory.

Edwards, D. J. & Silverberg, L. M. (2010). Autonomous Soaring: The Montague Cross-Country Challenge. *Journal of Aircraft*, *47*(5).

Fonseka, E. (2007). *Modeling and Flying Thermal Tubes with an UAV* (Master thesis, ETH Zurich).

Gordon, N., Salmond, D., & Smith, A. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F*, *140*(2), 107–113.

Hazard, M. W. (2010). Unscented Kalman Filter for Thermal Parameter Identification. In *48th AIAA Aerospace Sciences Meeting*.

Hinz, S. (2005). Fast and subpixel precise blob detection and attribution. In *IEEE International Conference on Image Processing 2005* (Vol. 3, pp. III–457).

Kümmerle, J., Hinzmann, T., Siegwart, R., & Gilitschenski, I. (2016). Real-time detection and tracking of multiple humans from high birds-eye views in the visual and infrared spectrum. In *Advances in Visual Computing - 12th International Symposium (ISVC), Las Vegas, NV, USA*.

Leutenegger, S., Melzer, A., Alexis, K., & Siegwart, R. (2014). Robust State Estimation for Small Unmanned Airplanes. In *IEEE Conference on Control Applications*.

Li, Z. M. (2010). *Modeling and Simulation of Autonomous Thermal Soaring with Horizon Simulation Framework* (Master thesis, California Polytechnic State University).

Malaver, A. J. R., Gonzalez, L. F., Motta, N., & Villa, T. F. (2015). Design and Flight Testing of an Integrated Solar Powered UAV and WSN for Remote Gas Sensing. In *IEEE Aerospace Conference*.

Meier, L., Honegger, D., & Pollefeys, M. (2015). PX4: A Node-Based Multithreaded Open Source Robotics Framework for Deeply Embedded Platforms. In *IEEE International Conference on Robotics and Automation (ICRA)*.

Nikolic, J., Rehder, J., Burri, M., Gohl, P., Leutenegger, S., Furgale, P. T., & Siegwart, R. Y. (2014). A Synchronized Visual-Inertial Sensor System with FPGA Pre-Processing for Accurate Real-Time SLAM. In *IEEE International Conference on Robotics and Automation (ICRA)*.

Noth, A. (2008). *Design of Solar Powered Airplanes for Continuous Flight* (Doctoral dissertation, ETH Zurich).

Oettershagen, P. (2017). *High-Fidelity Solar Power Income Modeling for Solar-Electric UAVs: Development and Flight Test Based Verification*. ETH Zurich. Available at http://arxiv.org/abs/1703.07385.

Oettershagen, P., Melzer, A., Mantel, T., Rudin, K., Lotz, R., Siebenmann, D., ... Siegwart, R. (2015). A Solar-Powered Hand-Launchable UAV for Low-Altitude Multi-Day Continuous Flight. In *IEEE International Conference on Robotics and Automation (ICRA)*.

Oettershagen, P., Melzer, A., Mantel, T., Rudin, K., Stastny, T., Wawrzacz, B., ... Siegwart, R. (2016). Perpetual flight with a small solar-powered UAV: Flight results, performance analysis and model validation. In *IEEE Aerospace Conference*.

Oettershagen, P., Melzer, A., Mantel, T., Rudin, K., Stastny, T., Wawrzacz, B., ... Siegwart, R. (2017). Design of small hand-launched solar-powered UAVs: From concept study to a multi-day world endurance record flight. *Journal of Field Robotics (JFR)*, *34*(7), 1352–1377.

Oettershagen, P., Stastny, T., Mantel, T., Melzer, A., Rudin, K., Gohl, P., ... Siegwart, R. (2015). Long-Endurance Sensing and Mapping using a Hand-Launchable Solar-Powered UAV. In *Field and Service Robotics (FSR)*.

Pamadi, B. N. (2004). *Performance, Stability, Dynamics, and Control of Airplanes* (2nd ed.). AIAA Education Series.

Park, S., Deyst, J., & How, J. P. (2007). Performance and Lyapunov stability of a nonlinear path following guidance method. *Journal of Guidance, Control, and Dynamics*, *30*(6), 1718–1728.

Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., ... Ng, A. Y. (2009). ROS: an open-source Robot Operating System. In *ICRA workshop on open source software* (3.2, p. 5). 3. Kobe, Japan.

Raymer, D. P. (2006). *Aircraft Design: A Conceptual Approach* (4th ed.). AIAA Education Series.

Reichmann, H. (1988). *Cross-Country Soaring*. Thomson Publications.

Suzuki, S. et al. (1985). Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, *30*(1), 32–46.

Tang, Y., Wu, Y., Wu, M., Wu, W., Hu, X., & Shen, L. (2009). INS/GPS Integration: Global ObservabilityAnalysis. *IEEE Transactions on Vehicular Technology*, *58*(3).

Tischler, M. B. & Remple, R. K. (2012). *Aircraft and rotorcraft system identification: engineering methods with flight-test examples*. American Institute of Aeronautics and Astronautics (AIAA).

U.S. Federal Aviation Administration. (2013). *Glider Flying Handbook*. U.S. Department of Transportation. Retrieved from https://www.faa.gov/regulations_policies/handbooks_manuals/aircraft/glider_handbook/

Vempati, A., Agamennoni, G., Stastny, T., & Siegwart, R. (2015). Victim Detection from a Fixed-Wing UAV: Experimental Results. In *International Symposium on Visual Computing (ISVC)* (pp. 432–443). Springer.

Wan, E. & Merwe, R. V. D. (2001). The Unscented Kalman Filter. In *Kalman Filtering and Neural Networks* (Chap. 7, pp. 221–280). Wiley.

Weider, A., Levy, H., Regev, I., Ankri, L., Goldenberg, T., Ehrlich, Y., ... Cohen, M. (2007). *SunSailor: Solar powered UAV*. Aerospace Engineering Faculty, Technion. Retrieved from http://webee.technion.ac.il/people/maxcohen/SunSailorArt19nov06.pdf

West, M. & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models* (2nd ed.). Springer.

Wharington, J. (1998). *Autonomous Control of Soaring Aircraft by Reinforcement Learning* (Doctoral dissertation, Royal Melbourne Institute of Technology).

Wu, Z., Fuller, N., Theriault, D., & Betke, M. (2014). A thermal infrared video benchmark for visual analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (pp. 201–208).