License Plate Recognition

What You Need to Know

800.434.1502 | T2systems.com

© Copyright 2015 T2 Systems, Inc. All Rights Reserved.

License Plate Recognition

Introduction

This document will provide an introduction to License Plate Recognition (LPR) and how it is used in the parking industry. Readers will understand important image considerations, learn how the LPR system works, and discover why the US has been described as the most difficult country for using LPR reliably.

What is LPR (ANPR)?

Extracting the license plate from an image into computerized text is necessary for systems processing where the license plate is a key field. Common applications for LPR include policing, tolling, parking enforcement, and identifying suspect vehicles.

LPR includes applications and methodologies for such a conversion. LPR includes elements of OCR (optical character recognition) which is used to convert the text from an image (like a scanned book or newspaper) into electronic format.

ANPR (Automatic Number Plate Recognition) is a synonym for LPR. ANPR is the term more frequently used in Europe where LPR has been around for several more years, and is more accurate than in North America.

LPR applications are generally categorized into three different types:

- Handheld LPR involves a handheld device, primarily for enforcement purposes, and assisting the officer in data entry of the license plate.
- Fixed LPR is the term used for a camera installed in a (semi)permanent location and tracking vehicles as they drive by. In parking this is most commonly found at the entry or egress of a parking facility, particularly airport facilities.
- Mobile LPR involves fixed cameras mounted to moving vehicle for enforcement, timing, and hotlist applications.

Image Considerations

What is the Role of the Camera?

A camera generates the image – or images – used by the LPR engine. The camera can provide both video or photo images, though video is often processed in a frame-by-frame manner.

The quality of the image is critical in improving read-accuracy. Most standalone LPR software packages can accept an image from different cameras and different sources. The image quality is more important than the camera; i.e. a single good image from a poor camera will work better than a poor image from a good camera.

The camera's ability to capture an image directly affects image quality. Critical factors include the camera's resolution (pixels), zoom (optical vs. digital), focal length (blur), etc. For successful LPR the best available image should be used; this usually means that the camera should be selected based on the needs of image quality. (Camera prices range from several hundred to several thousand dollars depending on feature set.)

Enviromental Conditions

Environmental conditions affect the image quality. Lighting is the most important factor for capturing a good image; artificial lighting (constant or flash) is commonly employed in conjunction with LPR.

Other factors affecting the image quality include sunlight, moonlight, fog/rain, shadows, glare, and so forth. Some of these can be controlled or adapted for, others will rely on LPR software to compensate.

IR (Infrared) vs. White Light

Many LPR cameras are capable of taking infrared (IR) images as well as normal (white light) images. Depending on conditions, such as poor lighting, an IR image may be better for LPR than white light.

Because LPR success and confidence is directly related to the quality of the image IR (infrared) cameras take photos outside of the spectrum of white light; these images show a greater contrast of IR reflective materials, such as the coating on most modern license plates.

The photos below are taken with a hybrid camera, one lens recording white light (left) and one recording IR (right). While the quality of the plate in the picture is not significantly different, it is much easier for the LPR engine to find the plate within the image.

There is some debate as to which image type – white light or IR – is better for LPR. Both have advantages and disadvantages. For example the road salt used for snow melting often splatters onto license plates. In a white light image the salt specs are invisible, but in an IR image they appear black.

Stand-Alone, Integrated, and Online LPR

In a stand-alone LPR system the image source (i.e. camera) is independent from the LPR engine, and often there are different manufacturers for the camera and the engine. Cost and configurability are advantages of this type of setup; complexity and compatibility are drawbacks.

In an integrated LPR system both the camera and the LPR engine are integrated, and provided by the same manufacturer. Often the LPR engine runs on a logic board or industrial PC inside the camera housing. External inputs (triggers) and outputs (network) are also available in this configuration. This type of setup is usually simpler but more expensive.

Finally, online ("cloud-based") LPR uses back-end processing, usually over the Internet, for camera sources that are far away. Advantages of this configuration include the ability for high-end processing and multiple camera feeds; disadvantages include the additional cost and time lag due to image transmission.

How LPR Works

Inside the System

The LPR "engine" is a software package; the primary input is an image file containing a license plate, and the primary output is the text of the plate. There are various other inputs to assist the LPR engine in decoding, and there are various other outputs as a result of the decoding.

The LPR engine receives the image in one of two ways: either the image is captured by an LPR camera and fed to the engine (integrated system, usually real-time), or alternatively the image can be delivered electronically to the LPR engine (remote processing, sometimes real-time and sometimes batch).

Some LPR engines will also accept "hints" to assist with decoding. These hints include state/region designations (to help with color detection), letter sizing (for fixed distance plates), or plate angle. These secondary inputs vary by manufacturer.

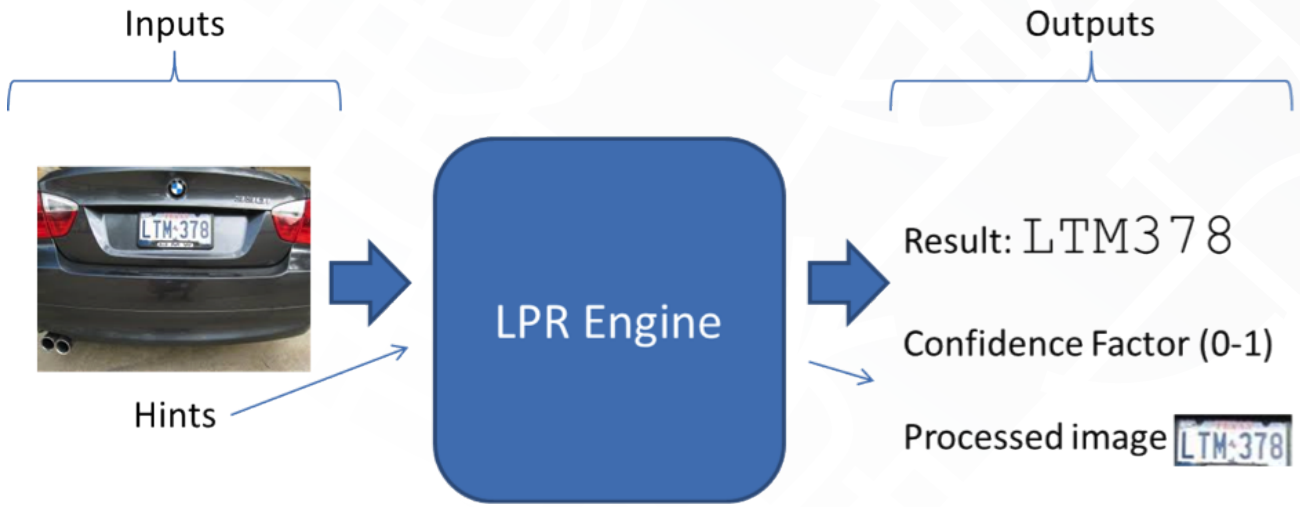

1. Yellow rectangles mark the areas that could be a license plate

2. Letters may be different sizes due to skew

3. Filters help isolate contrast; in this case the shadow along the baseline of the letters

4 . Whitespace is assumed to be vertical between letters, so kerning should be avoided

Generally speaking there are six steps involved in most plate recognition algorithms, as listed here:

Step 1: Localization or Framing

The LPR engine attempts to identify the position of the license plate within the image. This is called localization or framing. In this manner the LPR engine can focus on the plate and disregard any extraneous data in the picture.

Some of the LPR engines will also look at information outside the license plate frame and report additional information like the vehicle make or model (based on the logo and other text on the vehicle) as well as the vehicle color (sampled from the vehicle body).

Step 2: Orientation and Sizing

Depending on the angle of the license plate there may be a need for special adjustments needed to decode the plate. These mathematical, angular corrections assist in decoding plates that are taken from above (overhead camera), or from the side of a plate (parallel parked), and they correct for perspective and rotation.

Step 3: Normalization

Shadows and shading can present a problem for LPR, so various filters are used to reduce or eliminate these elements. Edge detection is a common filtering technique for license plates, and works well when there is high contrast between the text and background.

In more sophisticated systems multiple images of the same plate are blended together to create a clearer, more defined image for the processing engine. (This is similar to the HDR feature found on some electronic cameras and smartphones). Other image clarification techniques include applying various filters to make the image more readable and then using elements of each filtered image.

Step 4: Character Segmentation

This involves identifying the individual characters of the license plate using various techniques. Most commonly, this includes whitespace delineation (finding the spaces between the letters).

In a fixed space font the characters are easier to separate because the spacing pattern is predictable. When spacing is variable the LPR engine is more prone to error.

In some regions the plate manufacturer assists with LPR by adding kerning (adjusting the amount of white space) between characters that might be hard to decipher when letters are too close (for example VV can look like W, but kerning can help separate the two Vs).

5. Do you think this driver of plate "8B88B8B" is trying to thwart the LPR engine?

6. Zombie Apocalypse? Personalized plates often break the rules for character patterns.

Step 5: Optical Character Recognition

When the individual characters are separated, they are then recognized by various Optical Character Recognition (OCR) techniques. Pattern matching, edge tracing, proportion and pixel repetition are common methods for recognizing characters.

Using guesswork for matching contributes to the idea of a "confidence factor", which many LPR providers offer as an output of the LPR process. Overall confidence in a successful read is a factor of the individual character recognition as well as matching patterns suggested in the hints or tuning of the engine.

Some LPR engines will provide multiple options for a recognized plate; some letters and digits are virtually identical (I versus 1, O versus 0, B versus 8). The next step can sometimes aid in determining which option is more appropriate

Step 6: Syntactical/Geometric Analysis

This final step checks the characters and sequence against regionally specific rules.

In some cases this could be a list of known license plates (a hotlist or a whitelist, for example); this provides much higher confidence in plate recognition.

The introduction of personalized plates makes regional analysis more difficult since these patterns often break the standardized rules.

Confidence Factor

A confidence factor can be generated by the LPR engine as it processes the image of the license plate. It is an optional output of most LPR engines. Each decision (or guess) that the LPR engine makes results in some possibility of error – and the sum of these errors relates to the confidence in the decoding (fewer errors means higher confidence). Based on software tuning for specific plates the LPR engine is able to recognize when some guesses are better than others.

Confidence in a successful read is primarily related to recognition of the individual characters. However also factoring into the confidence are the patterns of letters, matching to a hotlist, and the quality of the original image. Some LPR engines will provide multiple output strings with a confidence factor for each (for the purpose of matching against a wildcard list).

Confidence is usually represented as a decimal between 0 and 1. For example, 0.85 represents an 85% confidence in the resulting license plate character string matching the original image. When the confidence is high it may be acceptable to consider the output as properly decoded; when the confidence is low a secondary system (or human) may be needed.

Whitelist / Hotlist

The term "whitelist" or "hotlist" is used to describe a list of known license plates. Usually the whitelist refers to plates that are valid (allowed to park), while hotlist are plates that are not valid (stolen, scofflaw, etc).

Having a whitelist or hotlist makes LPR easier because it gives the engine a smaller universe of matching values. For example if the whitelist contains B0123 but the LPR returns 8OI23 then there is a high probability these are the same plate (because B / 8 are similar, 0 / O are similar, and 1 / I are similar).

Other Elements

Some LPR engines are also able to detect information in addition to the plate. Some will look at the marque (the manufacturer's logo) and determine the vehicle make based on this information.

Vehicle model may also be available using various techniques like the shape/outline of the vehicle, or simply by reading the name of the model from the rear of the vehicle. Additionally, vehicle color is sometimes detectable, though lighting conditions play a large factor here.

LPR in North America

Additional Challenges with North American Plates

LPR works better in Europe because there is greater standardization of the license plates. Success rates of 98% or more are common using European plates, but success rates drop dramatically (frequently to below 80%) in the US and Canada for various reasons.

State Proximity

In the United States, and particularly near the intersection of multiple states, there is a significant overlap of plates from different origins. Along with the challenges of background and foreground color–the typeface (font) in different states will also vary. The LPR engine will have to work harder, and likely make more errors, when decoding multiple states.

In Europe, where many of the LPR (ANPR) companies got their start, there is less variability in plate color and typeface, even in different countries.

Pretty Backgrounds

States and provinces generate additional revenue by offering a level of personalization to plates (see also plate types below). Different plate backgrounds are one such offering.

LPR systems run into difficulty with varying backgrounds because the LPR engine needs to compensate for the background when reading the foreground letters. Backgrounds with low-contrast to the character color, as well as gradients, make for more difficult recognition.

Stacked Letters

Stacked letters have been phased out over the years, though some are still around. LPR systems vary by how they are recognized – pre-plate or post-plate, upper or lower.

Plate Types

There is an inherent assumption that most plates are unique; this is a reasonable assumption for most situations but occasionally the same character sequence will appear with different plate types.

Plate types are sometimes related to the type of vehicle (truck, trailer, taxi, etc) or special interest groups (firefighter, veteran, kids first, etc). Different states determine their own plate types.

It is rare to encounter the same plate with a different type, but it happens from time to time. Allowed use of personalized plates also increases the chance of finding a duplicate. Using a whitelist or hotlist can be helpful here, though most LPR will not differentiate by type. Subsequently there will be plate type exceptions for the same plate characters.

Summary and Conclusions

LPR is rapidly being accepted in the parking industry to augment (or sometimes replace) use of credentials, but LPR is not an exact science and successful recognition rates vary wildly. Purchasers and consumers of LPR technology should expect recognition rates under 90% and be pleasantly surprised at anything higher. In addition, North American consumers will see even more variability from locale-to-locale.