Webcam Clip Art:
Appearance and Illuminant Transfer from Time-lapse Sequences

Jean-François Lalonde, Alexei A. Efros,
and Srinivasa G. Narasimhan

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Snow in mountains?

Sunny beach?

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Most of us use webcams to plan our daily activities: for example,
- we want to go skiing, but want to make sure the mountain’s got plenty of snow;
- but the beach is sunny today so maybe we go there instead;
- and we can even see if we’re going to hit some traffic on our way.

But there’s actually a lot more to it in a webcam!

Snow in mountains?

Sunny beach?

Traffic?

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Most of us use webcams to plan our daily activities: for example,
- we want to go skiing, but want to make sure the mountain's got plenty of snow;
- but the beach is sunny today so maybe we go there instead;
- and we can even see if we're going to hit some traffic on our way.

But there's actually a lot more to it in a webcam!

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

It's this amazing thing where we can hold all the variables constant, except the illumination. So it's not just a bunch of pictures, but a scientific tool that we must use to understand appearance in the real world. In fact, there has been a number of people interested in webcams, but most of the interest so far has been in vision.
Today I'm going to show you how we can use the webcams of the world to provide a novel understanding of natural illumination in real images, and how that can be used in graphics. For example, a popular graphics application which requires a good understanding of illumination is object insertion.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

It’s this amazing thing where we can hold all the variables constant, except the illumination. So it’s not just a bunch of pictures, but a scientific tool that we must use to understand appearance in the real world. In fact, there has been a number of people interested in webcams, but most of the interest so far has been in vision.

Today I’m going to show you how we can use the webcams of the world to provide a novel understanding of natural illumination in real images, and how that can be used in graphics. For example, a popular graphics application which requires a good understanding of illumination is object insertion.

It’s this amazing thing where we can hold all the variables constant, except the illumination. So it’s not just a bunch of pictures, but a scientific tool that we must use to understand appearance in the real world. In fact, there has been a number of people interested in webcams, but most of the interest so far has been in vision.

Today I’m going to show you how we can use the webcams of the world to provide a novel understanding of natural illumination in real images, and how that can be used in graphics. For example, a popular graphics application which requires a good understanding of illumination is object insertion.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

It's this amazing thing where we can hold all the variables constant, except the illumination. So it's not just a bunch of pictures, but a scientific tool that we must use to understand appearance in the real world. In fact, there has been a number of people interested in webcams, but most of the interest so far has been in vision.
Today I'm going to show you how we can use the webcams of the world to provide a novel understanding of natural illumination in real images, and how that can be used in graphics. For example, a popular graphics application which requires a good understanding of illumination is object insertion.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Suppose city officials want to insert a Venus de Milo statue in that public square in Madrid, and visualize how it would look under different lighting conditions.

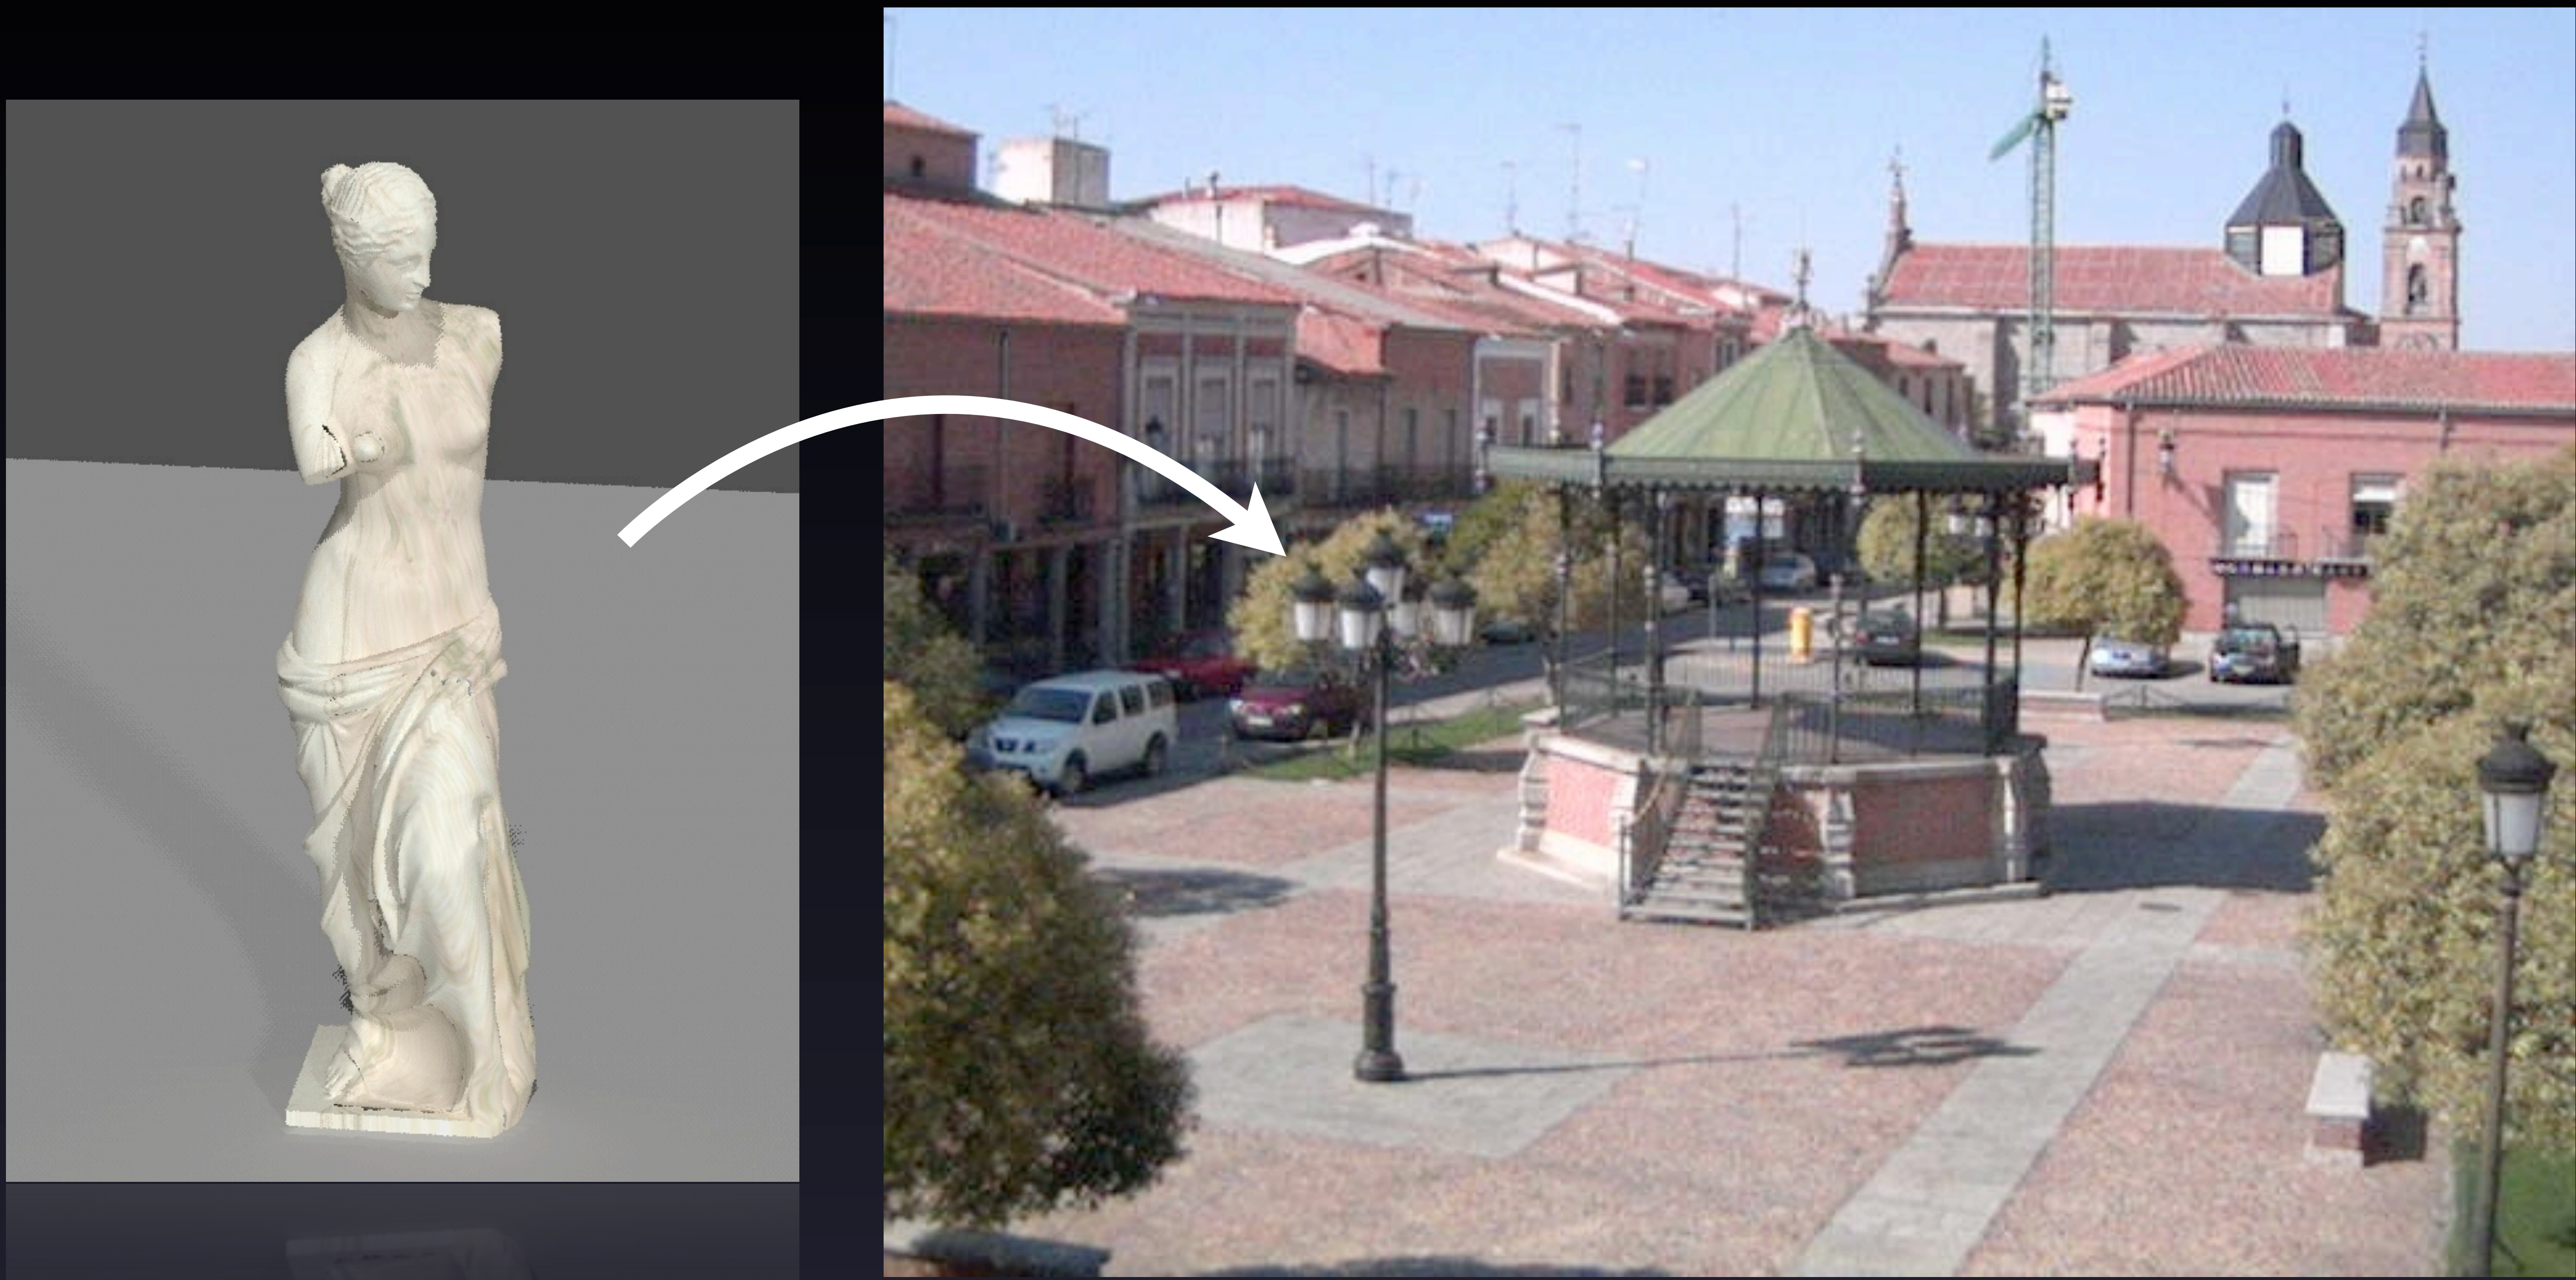

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Suppose city officials want to insert a Venus de Milo statue in that public square in Madrid, and visualize how it would look under different lighting conditions.

Measuring illumination [Debevec, 1998]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

5

We can use the method proposed by Debevec in 98, which directly measures the intensity of the light coming in from all directions by taking high-dynamic range photographs of a metallic sphere inserted into the scene. We can use this to realistically render our 3-D model. The downside of this method is that it requires a complex geometric and radiometric calibration procedure.

Measuring illumination [Debevec, 1998]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We can use the method proposed by Debevec in 98, which directly measures the intensity of the light coming in from all directions by taking high-dynamic range photographs of a metallic sphere inserted into the scene. We can use this to realistically render our 3-D model. The downside of this method is that it requires a complex geometric and radiometric calibration procedure.

Measuring illumination [Debevec, 1998]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We can use the method proposed by Debevec in 98, which directly measures the intensity of the light coming in from all directions by taking high-dynamic range photographs of a metallic sphere inserted into the scene. We can use this to realistically render our 3-D model. The downside of this method is that it requires a complex geometric and radiometric calibration procedure.

Measuring illumination [Debevec, 1998]

Complex calibration

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We can use the method proposed by Debevec in 98, which directly measures the intensity of the light coming in from all directions by taking high-dynamic range photographs of a metallic sphere inserted into the scene. We can use this to realistically render our 3-D model. The downside of this method is that it requires a complex geometric and radiometric calibration procedure.

Matching illumination [Lalonde et al., 2007]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Instead of measuring illumination, an alternative idea is the Photo Clip Art approach which we proposed 2 years ago: gather a lot of photos of statues, and look for one that has similar lighting conditions than our scene. Then all we have to do is to transfer its appearance onto the image, and we're done: no calibration or rendering is necessary. We've moved from measuring to matching illumination.

Unfortunately, the matching features we used were only able to represent rough illumination properties like sunny or overcast, and were not able to capture more complex effects like the sun direction. This can give us inconsistent results where shadows are pointing in the wrong direction as shown here.

Matching illumination [Lalonde et al., 2007]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Instead of measuring illumination, an alternative idea is the Photo Clip Art approach which we proposed 2 years ago: gather a lot of photos of statues, and look for one that has similar lighting conditions than our scene. Then all we have to do is to transfer its appearance onto the image, and we’re done: no calibration or rendering is necessary. We’ve moved from measuring to matching illumination.

Unfortunately, the matching features we used were only able to represent rough illumination properties like sunny or overcast, and were not able to capture more complex effects like the sun direction. This can give us inconsistent results where shadows are pointing in the wrong direction as shown here.

Matching illumination [Lalonde et al., 2007]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Instead of measuring illumination, an alternative idea is the Photo Clip Art approach which we proposed 2 years ago: gather a lot of photos of statues, and look for one that has similar lighting conditions than our scene. Then all we have to do is to transfer its appearance onto the image, and we're done: no calibration or rendering is necessary. We've moved from measuring to matching illumination.
Unfortunately, the matching features we used were only able to represent rough illumination properties like sunny or overcast, and were not able to capture more complex effects like the sun direction. This can give us inconsistent results where shadows are pointing in the wrong direction as shown here.

Matching illumination [Lalonde et al., 2007]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Instead of measuring illumination, an alternative idea is the Photo Clip Art approach which we proposed 2 years ago: gather a lot of photos of statues, and look for one that has similar lighting conditions than our scene. Then all we have to do is to transfer its appearance onto the image, and we're done: no calibration or rendering is necessary. We've moved from measuring to matching illumination.

Unfortunately, the matching features we used were only able to represent rough illumination properties like sunny or overcast, and were not able to capture more complex effects like the sun direction. This can give us inconsistent results where shadows are pointing in the wrong direction as shown here.

Matching illumination [Lalonde et al., 2007]

No sun direction

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Instead of measuring illumination, an alternative idea is the Photo Clip Art approach which we proposed 2 years ago: gather a lot of photos of statues, and look for one that has similar lighting conditions than our scene. Then all we have to do is to transfer its appearance onto the image, and we're done: no calibration or rendering is necessary. We've moved from measuring to matching illumination.

Unfortunately, the matching features we used were only able to represent rough illumination properties like sunny or overcast, and were not able to capture more complex effects like the sun direction. This can give us inconsistent results where shadows are pointing in the wrong direction as shown here.

Goal

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So our goal is to find a method that will capture the illumination conditions accurately in a physics-based fashion, without having to explicitly measure it. If we can do that, we'll be able to:
- create these metallic spheres without having to take photographs of them, which will allow us to realistically insert virtual 3-D objects into images with no calibration;
- accurately match illumination in object databases, allowing us to realistically transfer 2-D objects across images, capturing even the sun direction.
This is very hard to do in a single image, and while we have obtained preliminary results in our ICCV work this year, this was not good enough for graphics.

Goal

Capture illumination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So our goal is to find a method that will capture the illumination conditions accurately in a physics-based fashion, without having to explicitly measure it. If we can do that, we'll be able to:
- create these metallic spheres without having to take photographs of them, which will allow us to realistically insert virtual 3-D objects into images with no calibration;
- accurately match illumination in object databases, allowing us to realistically transfer 2-D objects across images, capturing even the sun direction.

This is very hard to do in a single image, and while we have obtained preliminary results in our ICCV work this year, this was not good enough for graphics.

Goal

Capture illumination without measuring it

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So our goal is to find a method that will capture the illumination conditions accurately in a physics-based fashion, without having to explicitly measure it. If we can do that, we'll be able to:
- create these metallic spheres without having to take photographs of them, which will allow us to realistically insert virtual 3-D objects into images with no calibration;
- accurately match illumination in object databases, allowing us to realistically transfer 2-D objects across images, capturing even the sun direction.

This is very hard to do in a single image, and while we have obtained preliminary results in our ICCV work this year, this was not good enough for graphics.

Goal

Capture illumination without measuring it

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So our goal is to find a method that will capture the illumination conditions accurately in a physics-based fashion, without having to explicitly measure it. If we can do that, we'll be able to:
- create these metallic spheres without having to take photographs of them, which will allow us to realistically insert virtual 3-D objects into images with no calibration;
- accurately match illumination in object databases, allowing us to realistically transfer 2-D objects across images, capturing even the sun direction.

This is very hard to do in a single image, and while we have obtained preliminary results in our ICCV work this year, this was not good enough for graphics.

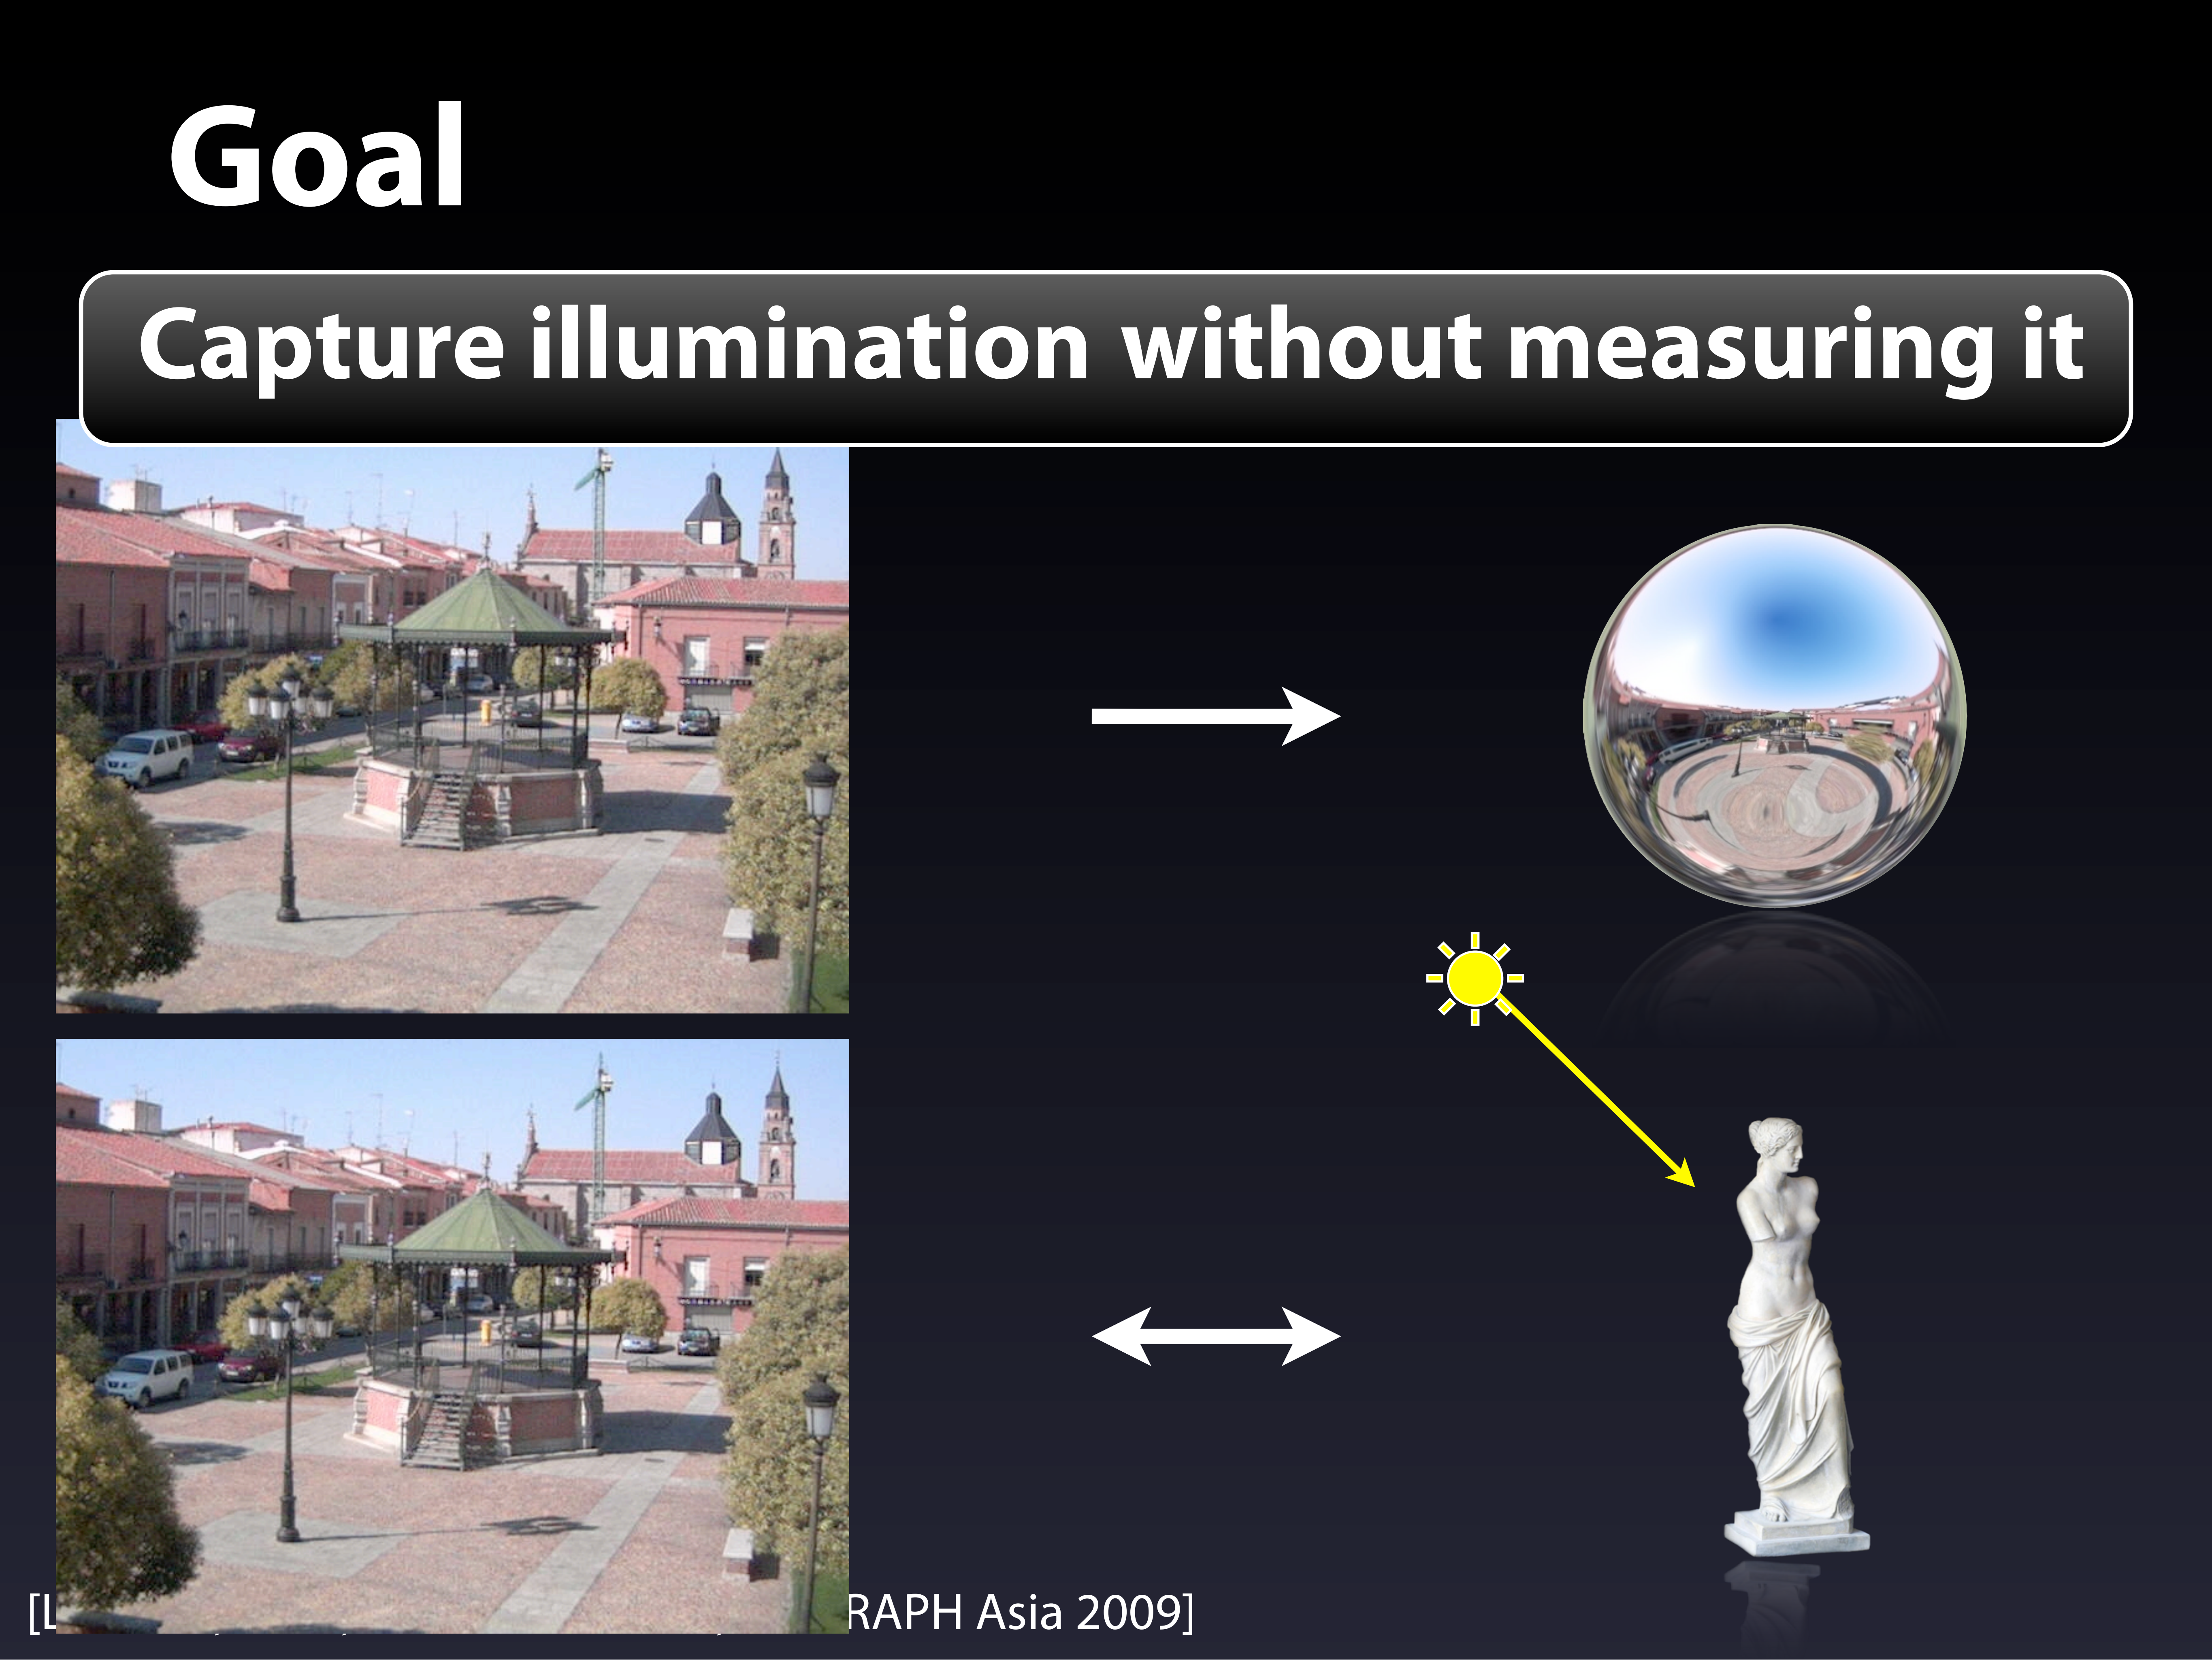

So our goal is to find a method that will capture the illumination conditions accurately in a physics-based fashion, without having to explicitly measure it. If we can do that, we'll be able to:
- create these metallic spheres without having to take photographs of them, which will allow us to realistically insert virtual 3-D objects into images with no calibration;
- accurately match illumination in object databases, allowing us to realistically transfer 2-D objects across images, capturing even the sun direction.

This is very hard to do in a single image, and while we have obtained preliminary results in our ICCV work this year, this was not good enough for graphics.

Our approach

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

But this is precisely where webcams can help us because they encode a lot more information about illumination! In this talk, I will show that we can achieve these goals and get high-quality results simply by capturing several photos of the scene taken over time. First, we will estimate the illumination conditions at each image in the sequence. From there, we will see how to generate synthetic metallic spheres, aka environment maps, that can be used to render a virtual 3-D object; and how we can define a matching measure suitable for 2-D object transfer. In both cases, this whole process is entirely automatic.

Our approach

Image sequence

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

But this is precisely where webcams can help us because they encode a lot more information about illumination! In this talk, I will show that we can achieve these goals and get high-quality results simply by capturing several photos of the scene taken over time. First, we will estimate the illumination conditions at each image in the sequence. From there, we will see how to generate synthetic metallic spheres, aka environment maps, that can be used to render a virtual 3-D object; and how we can define a matching measure suitable for 2-D object transfer. In both cases, this whole process is entirely automatic.

Our approach

Image sequence

⇨ Illumination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

But this is precisely where webcams can help us because they encode a lot more information about illumination! In this talk, I will show that we can achieve these goals and get high-quality results simply by capturing several photos of the scene taken over time. First, we will estimate the illumination conditions at each image in the sequence. From there, we will see how to generate synthetic metallic spheres, aka environment maps, that can be used to render a virtual 3-D object; and how we can define a matching measure suitable for 2-D object transfer. In both cases, this whole process is entirely automatic.

Our approach

Image sequence

Illumination

Environment maps

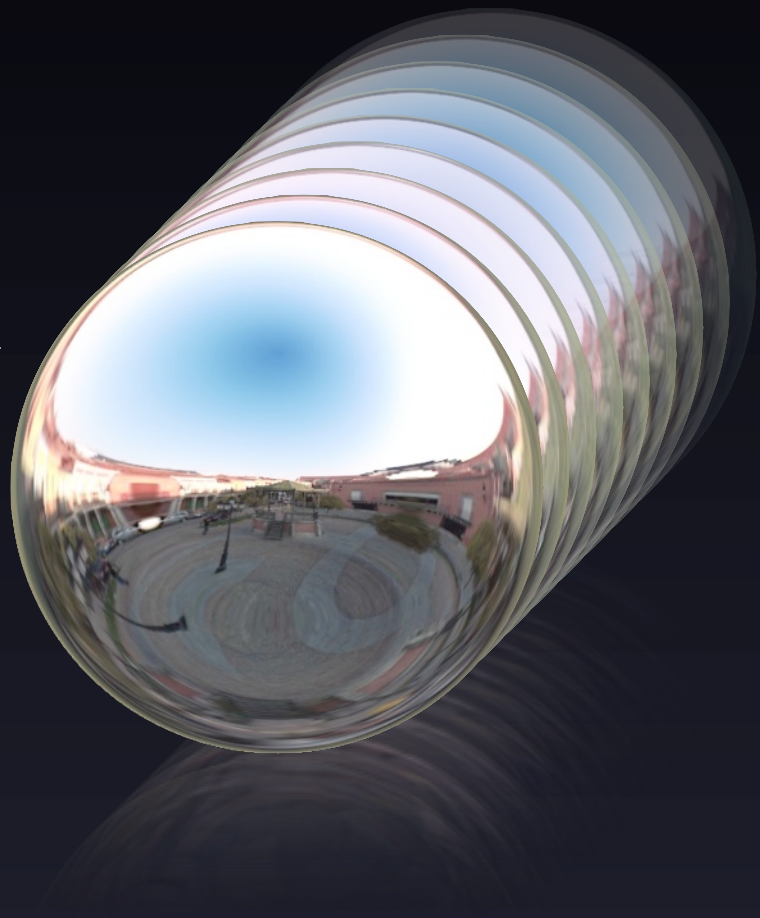

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

But this is precisely where webcams can help us because they encode a lot more information about illumination! In this talk, I will show that we can achieve these goals and get high-quality results simply by capturing several photos of the scene taken over time. First, we will estimate the illumination conditions at each image in the sequence. From there, we will see how to generate synthetic metallic spheres, aka environment maps, that can be used to render a virtual 3-D object; and how we can define a matching measure suitable for 2-D object transfer. In both cases, this whole process is entirely automatic.

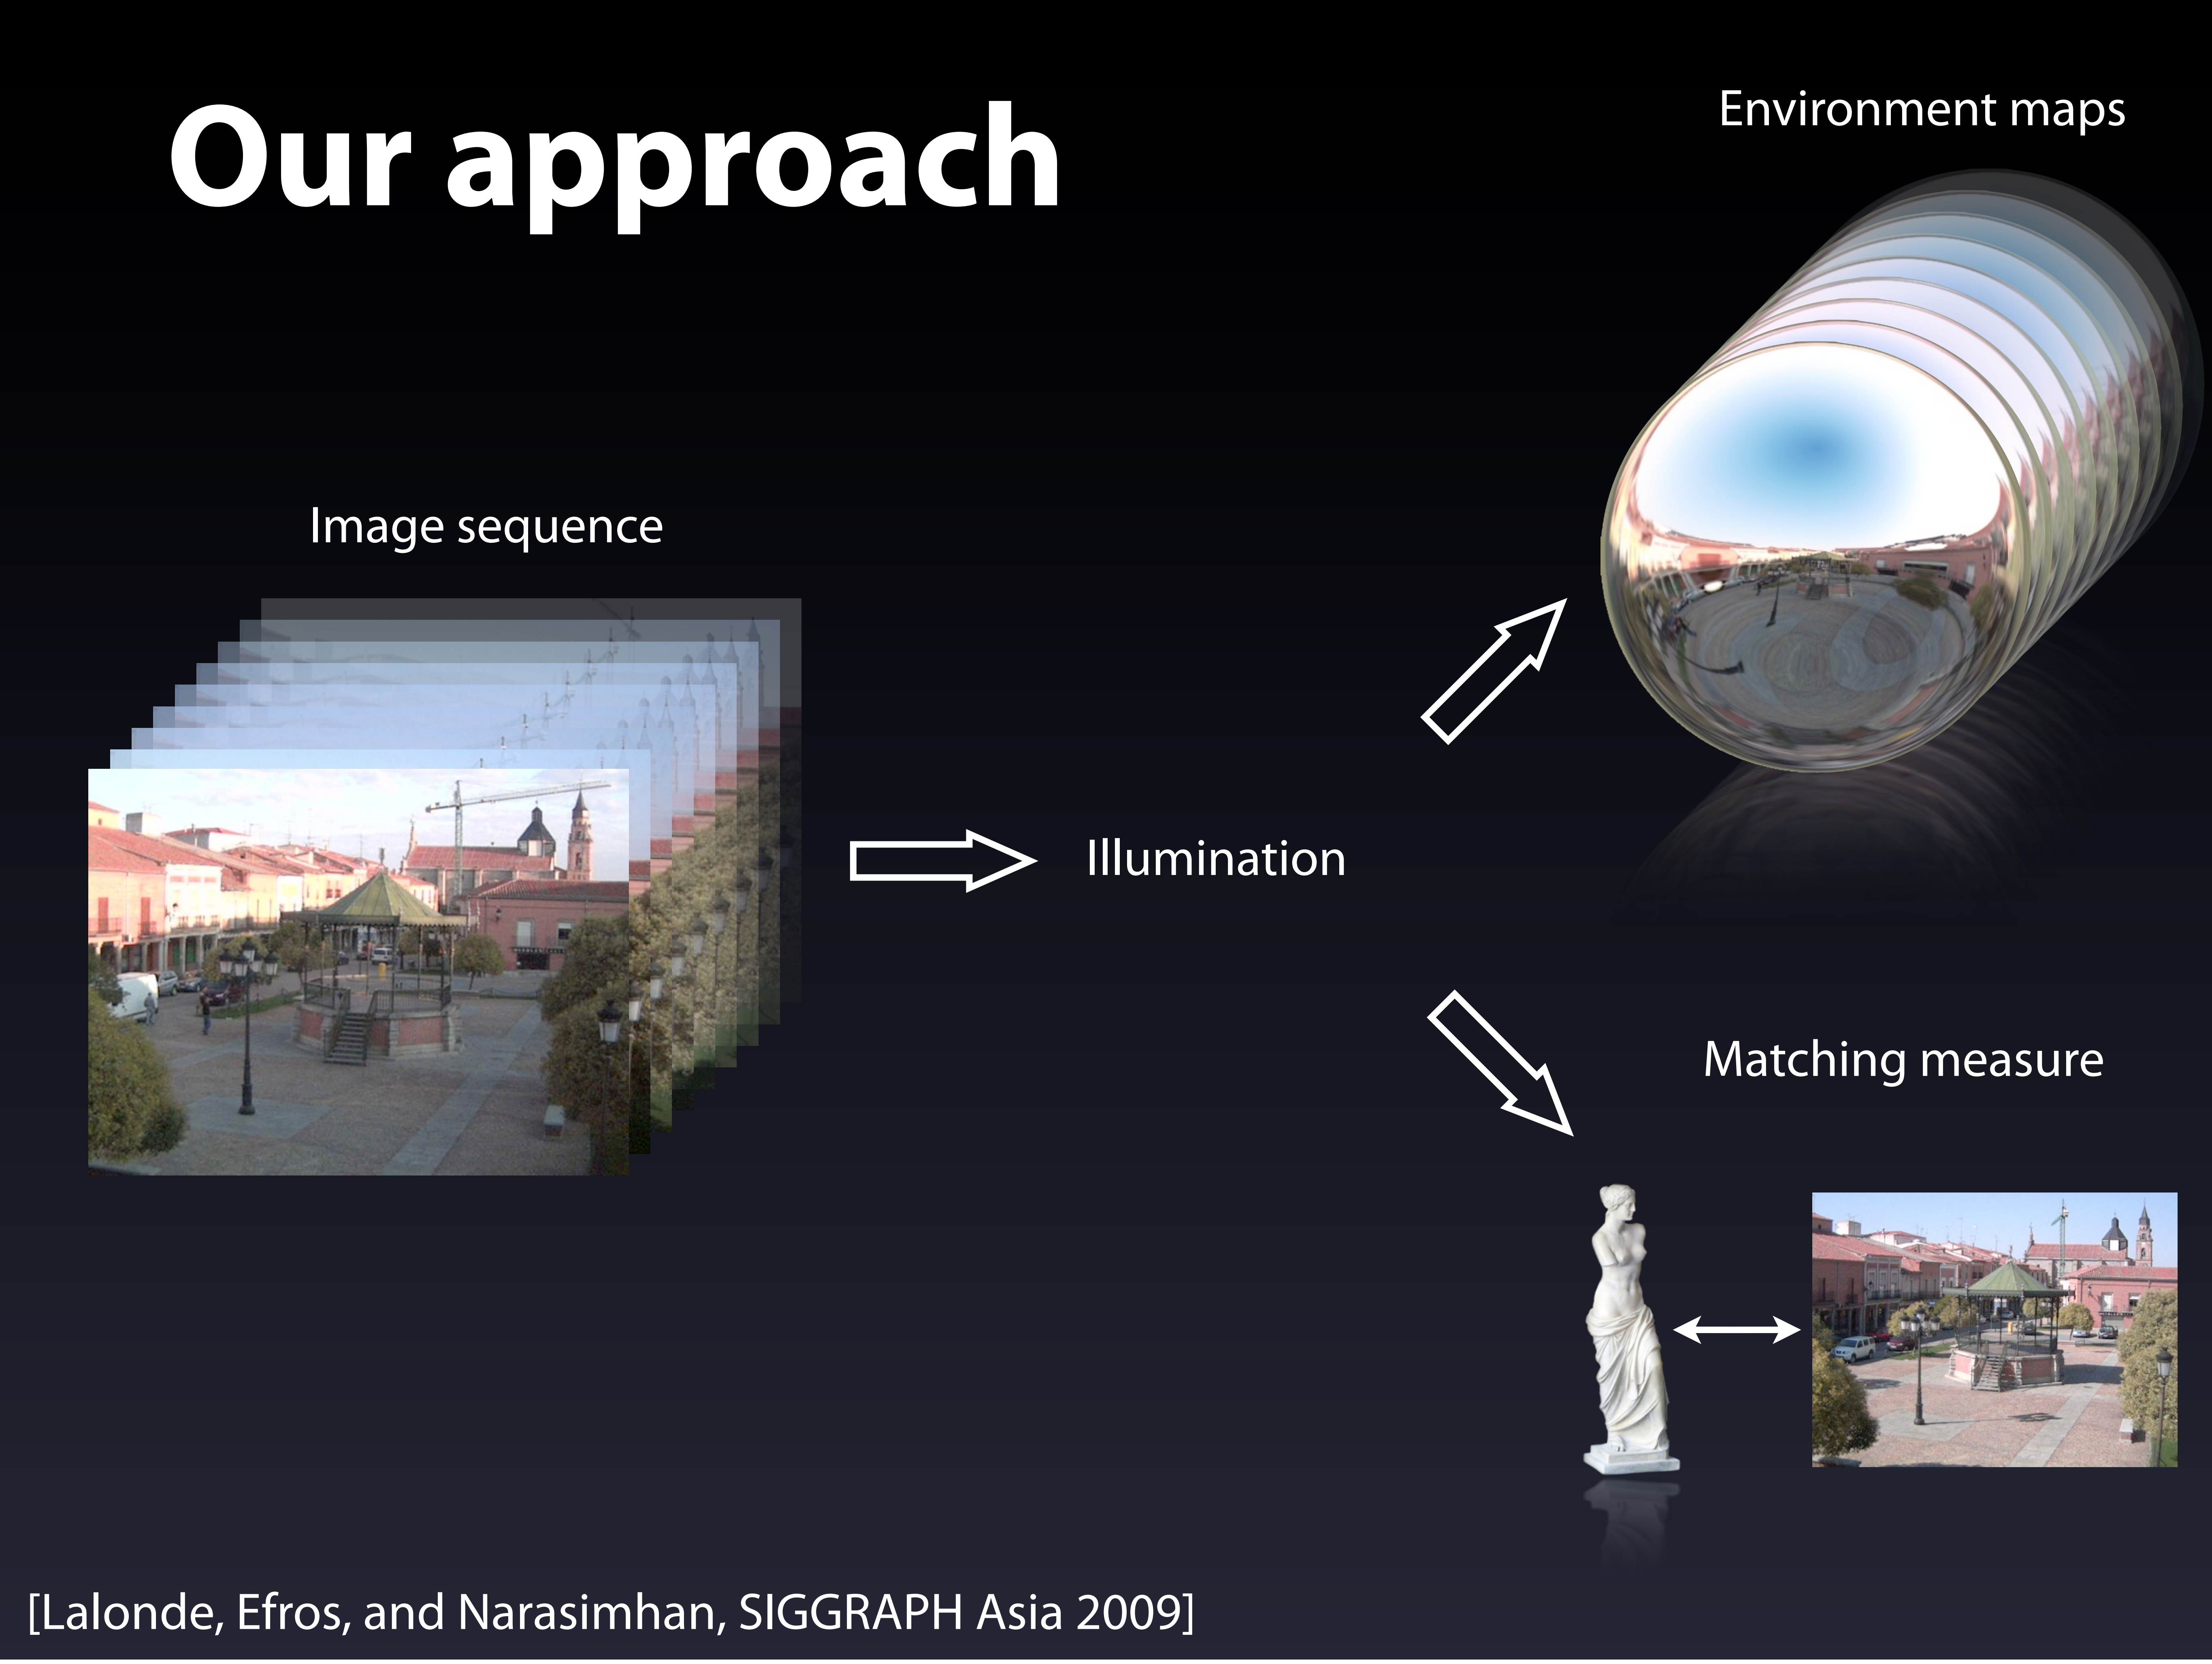

But this is precisely where webcams can help us because they encode a lot more information about illumination! In this talk, I will show that we can achieve these goals and get high-quality results simply by capturing several photos of the scene taken over time. First, we will estimate the illumination conditions at each image in the sequence. From there, we will see how to generate synthetic metallic spheres, aka environment maps, that can be used to render a virtual 3-D object; and how we can define a matching measure suitable for 2-D object transfer. In both cases, this whole process is entirely automatic.

Illumination from webcams

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Illumination from webcams

Image sequence

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Illumination from webcams

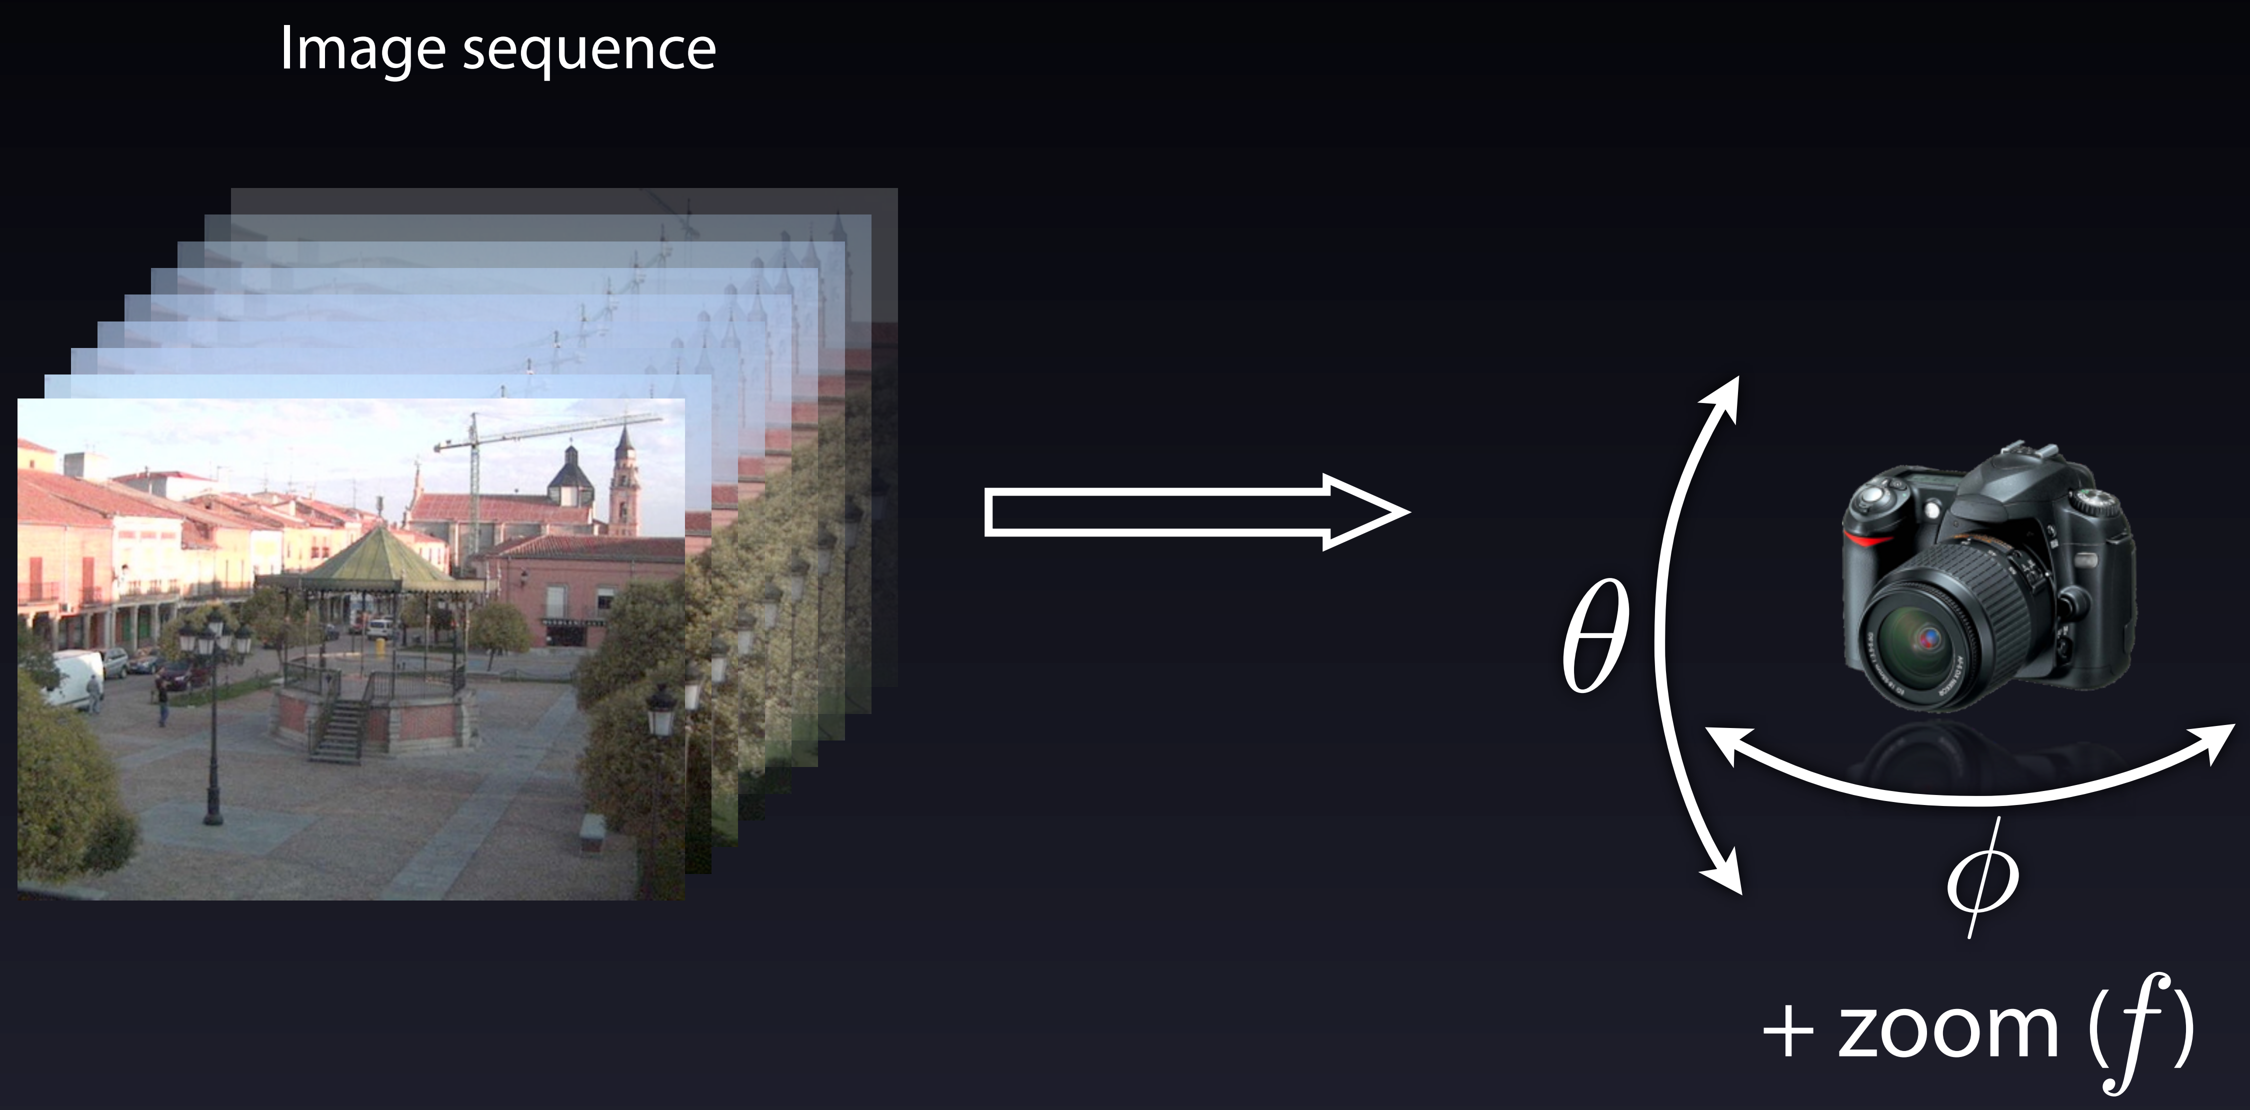

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Illumination from webcams

Image sequence

- sun position

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Illumination from webcams

Image sequence

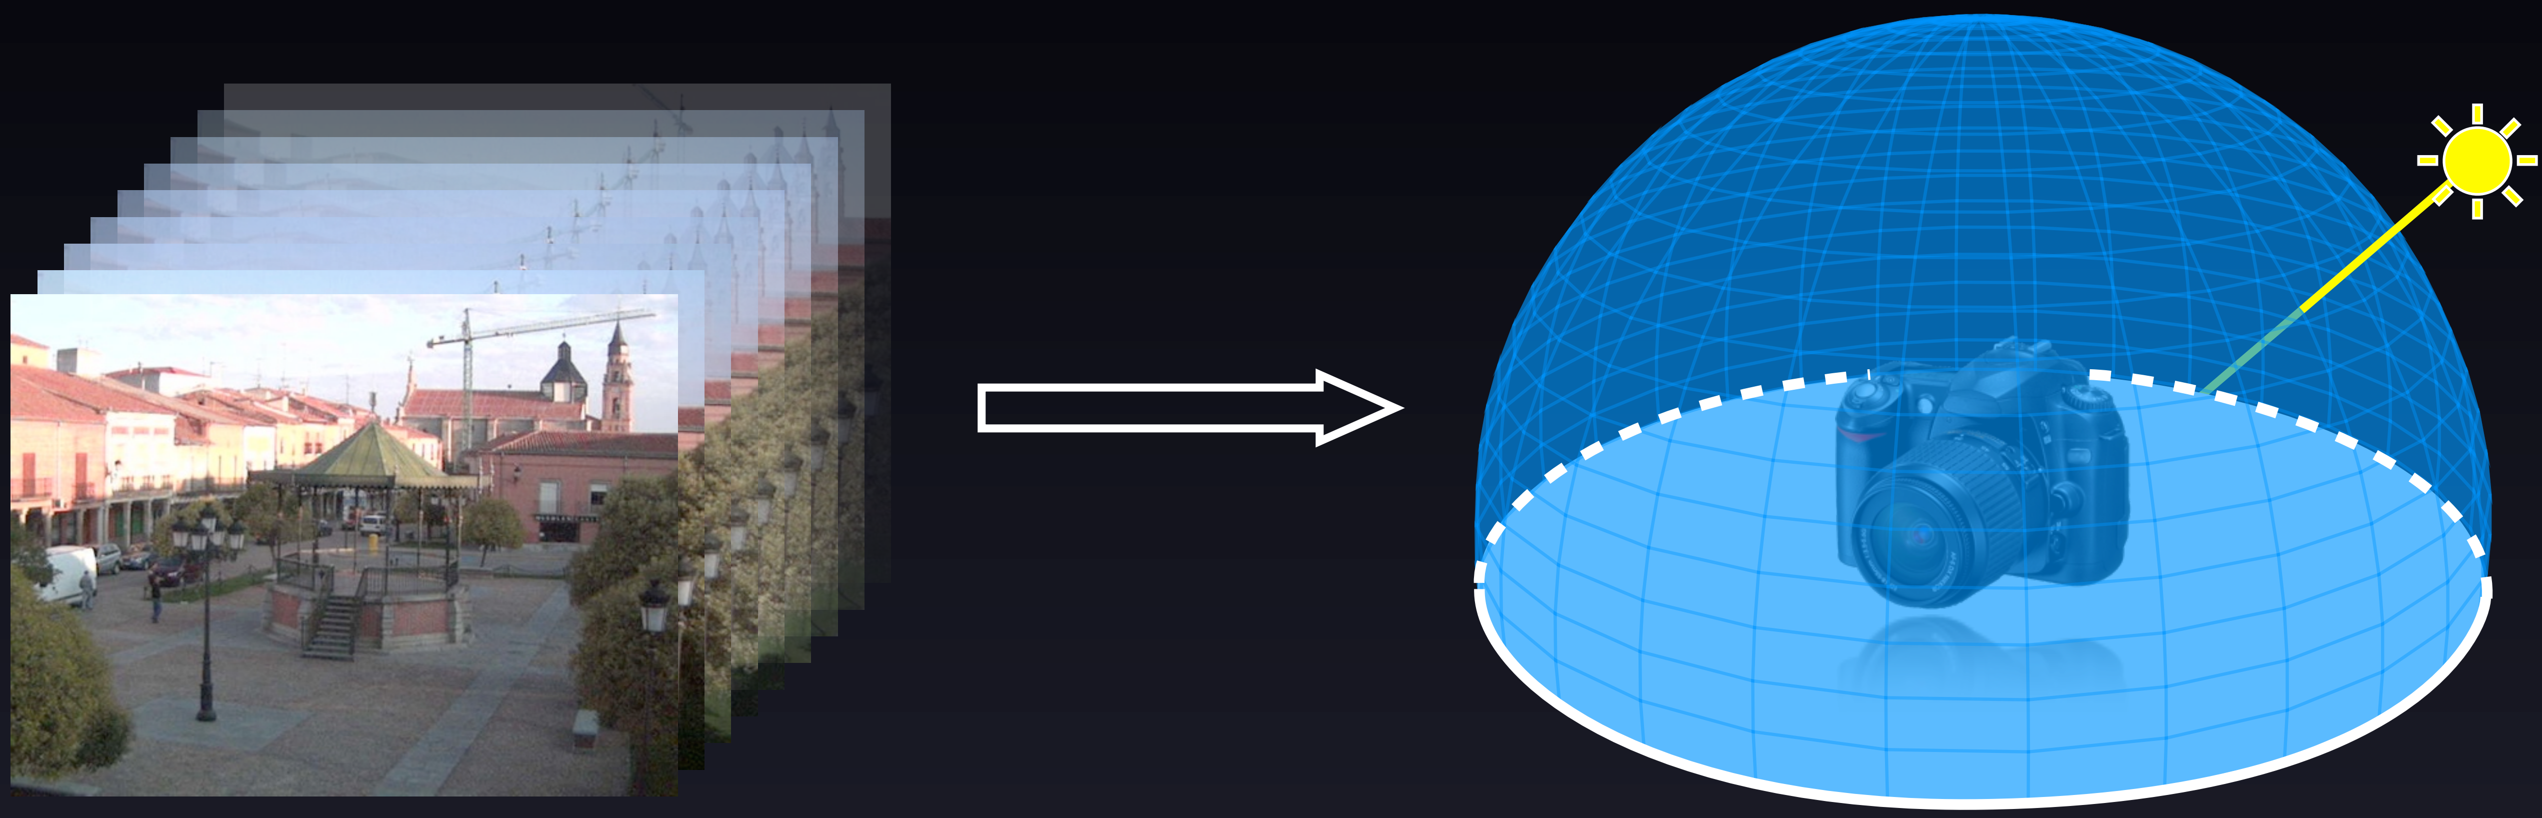

- sun position
- sky turbidity

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Illumination from webcams

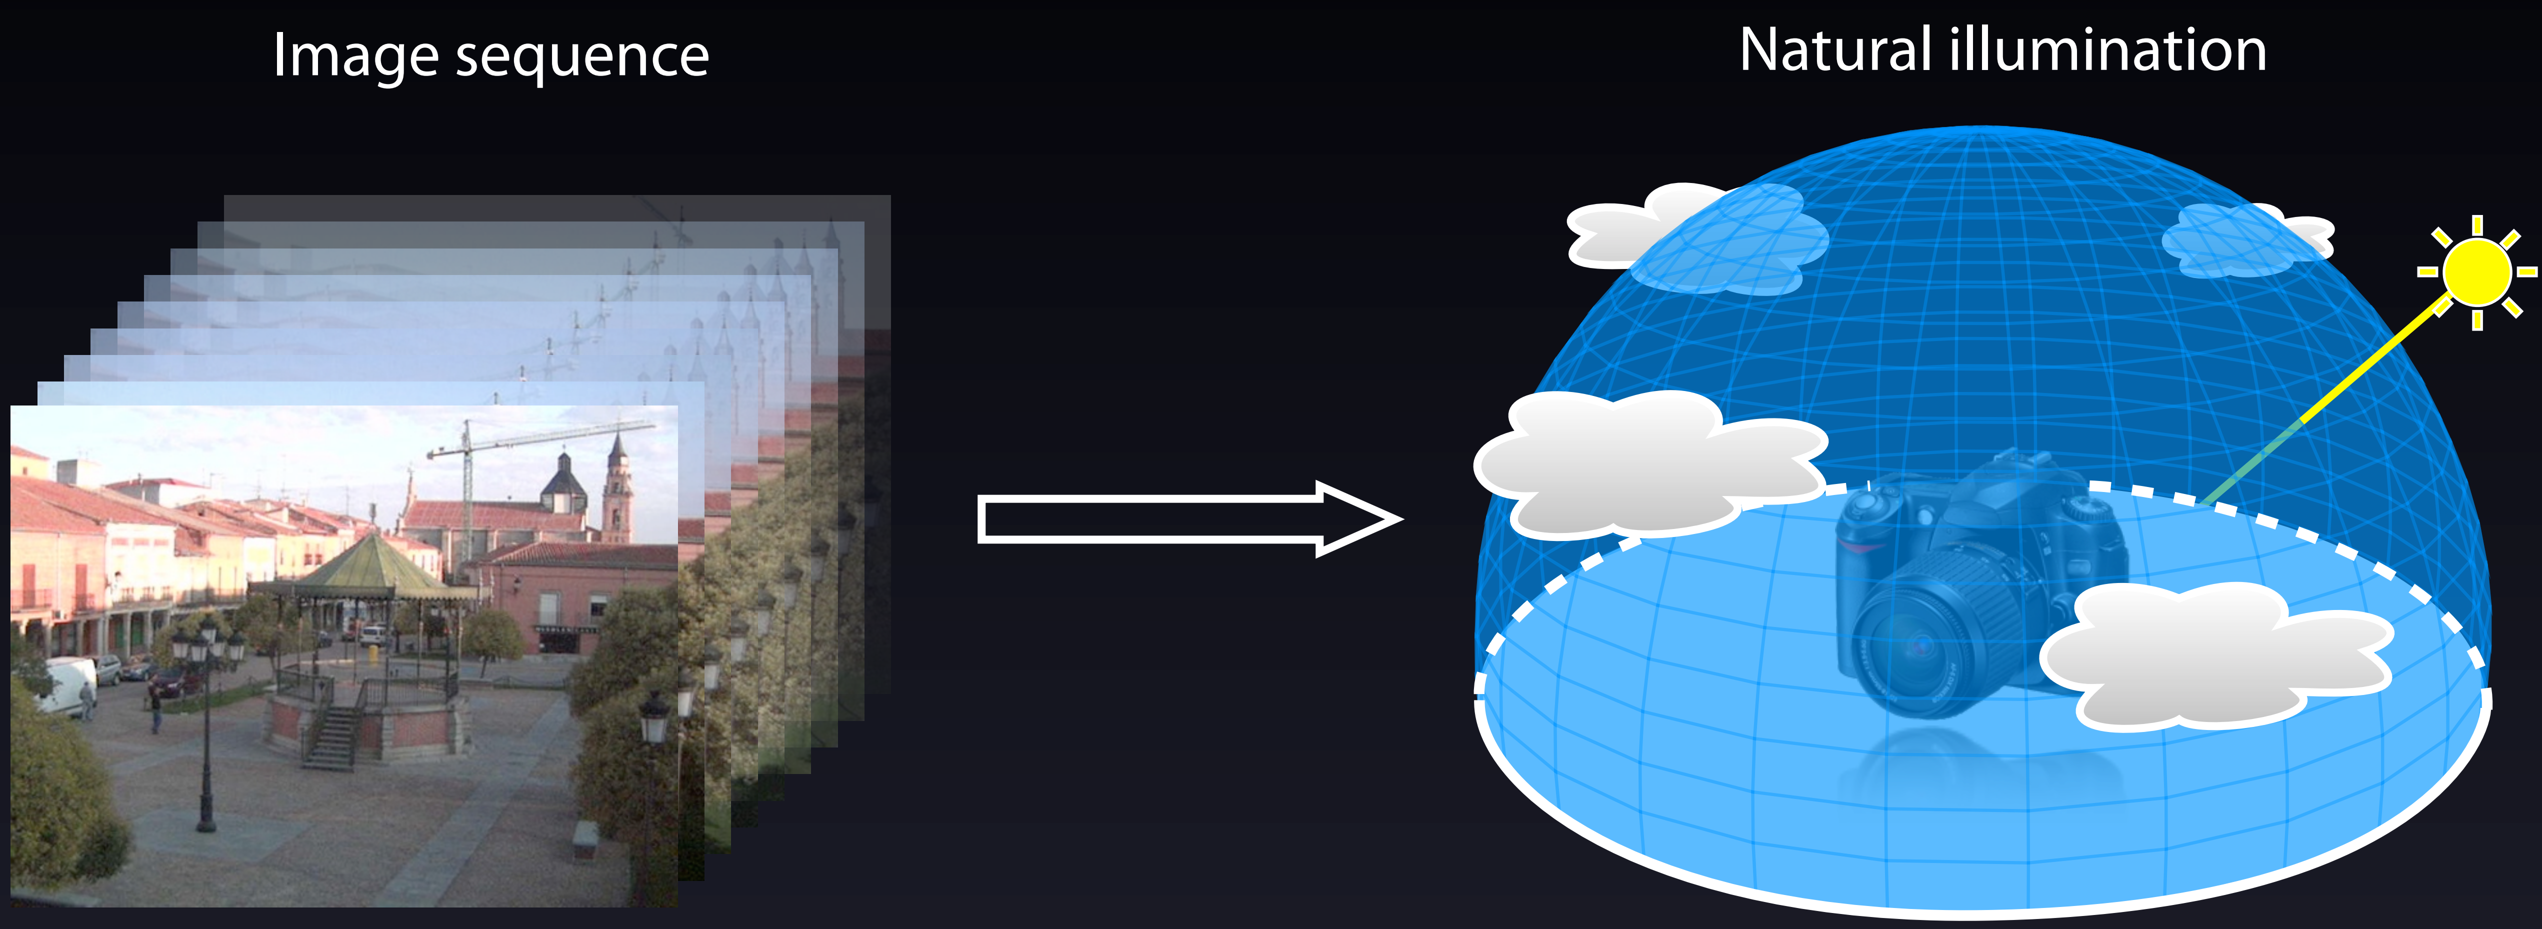

- sun position
- sky turbidity
- cloud segmentation

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Lalonde *et al.*, 2009]

The first step is to estimate the illumination conditions from the image sequence. For this we rely on our recent work, in which we showed how we can recover the camera orientation and zoom just by looking at the clear sky over time. Basically, the technique works by fitting the physically-based model of Perez et al. to several images of the clear sky taken at different times of day. Think about how useful that is: no matter how complicated the scene is, if we can see the sky even in a small portion of the image, we can use it to recover the camera parameters. Not only that: we can also use it to estimate the sun position relative to the camera, the sky turbidity, which is a measure of how clear it is, and even a cloud segmentation at each frame in the sequence.

Orientation & zoom from sky

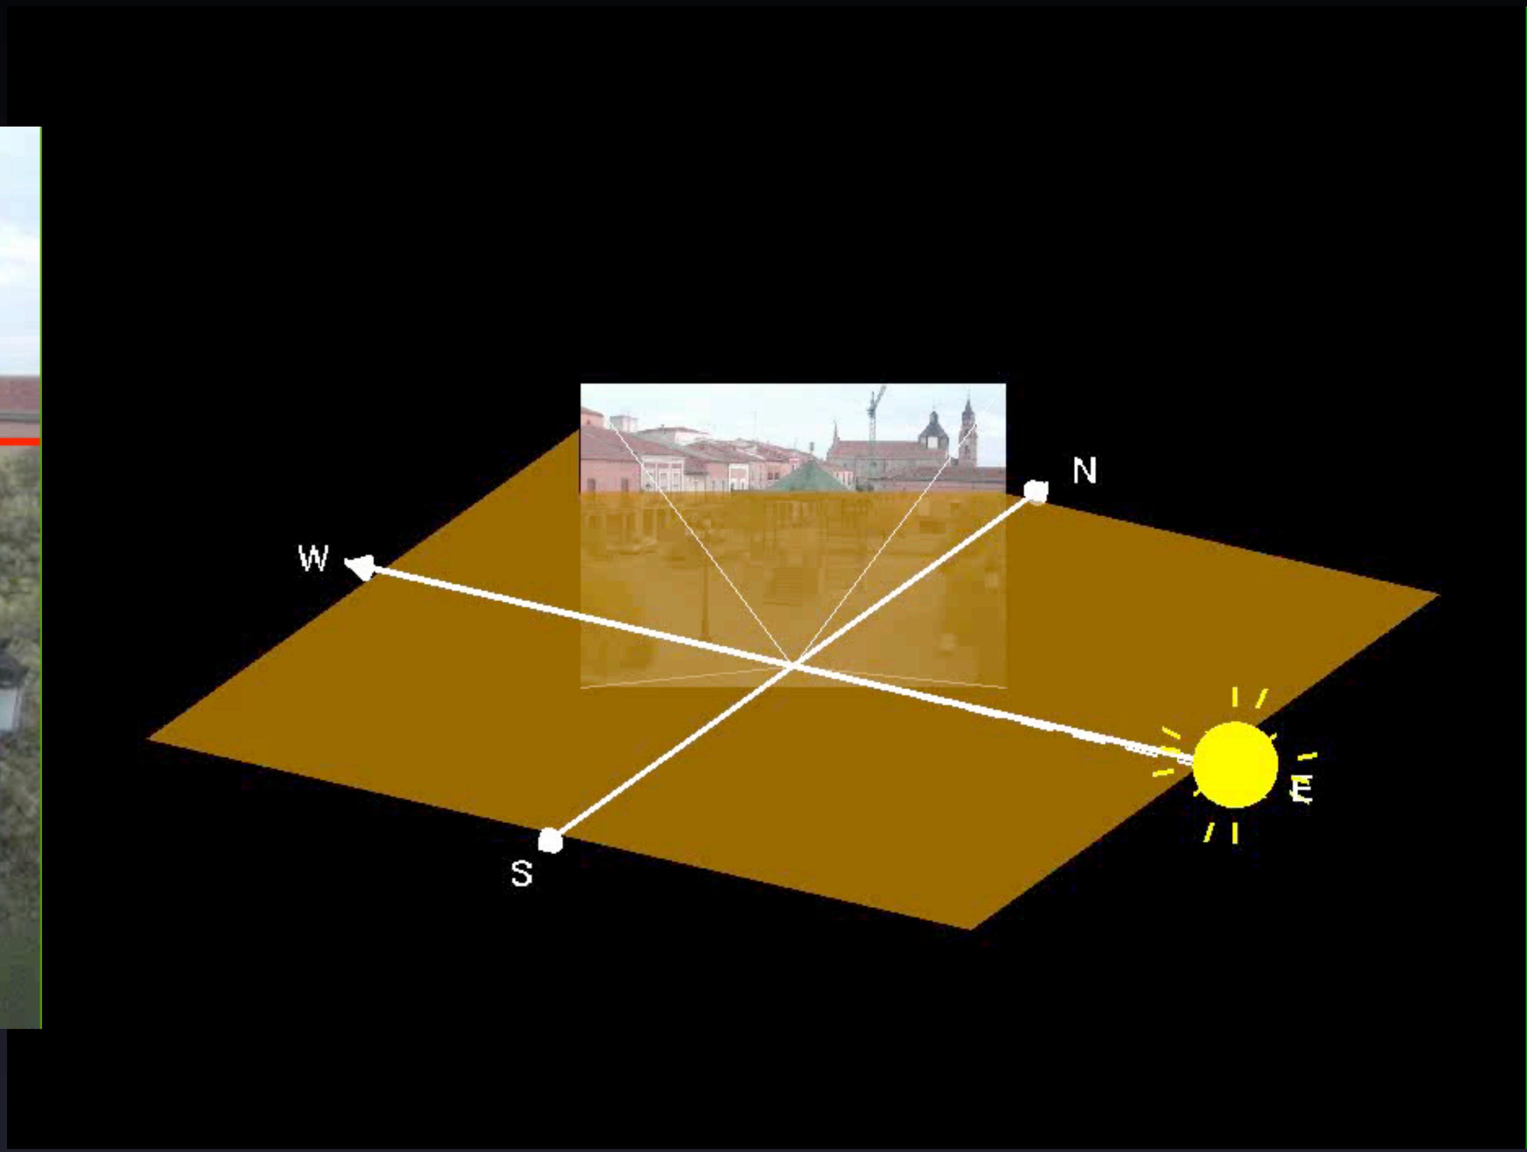

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Here is the result of applying this algorithm on a real image sequence. On the left is a single day from a webcam sequence, along with the estimated horizon line. On the right is a display showing the recovered camera orientation, zoom, and relative sun position at each frame. Look at how the shadows on the ground are consistent with the sun position.

Notice here that we can do all this even if the sun is never in the field of view of the camera, but we can still recover its position just by looking at the effect it has on the visible sky.

Is the sun shining on the scene?

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This is really nice, but unfortunately it doesn't capture the whole story: if we never see the sun, how can we know whether it's shining directly on the scene, or hiding behind a cloud? This is really important, because the scene will be lit very differently whether the sun is shining directly on it, or not. Unfortunately, the sky is not a reliable indicator. We have to look at the scene.

Is the sun shining on the scene?

Sun shining

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This is really nice, but unfortunately it doesn't capture the whole story: if we never see the sun, how can we know whether it's shining directly on the scene, or hiding behind a cloud? This is really important, because the scene will be lit very differently whether the sun is shining directly on it, or not. Unfortunately, the sky is not a reliable indicator. We have to look at the scene.

Is the sun shining on the scene?

Sun shining

Sun hiding

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This is really nice, but unfortunately it doesn't capture the whole story: if we never see the sun, how can we know whether it's shining directly on the scene, or hiding behind a cloud? This is really important, because the scene will be lit very differently whether the sun is shining directly on it, or not. Unfortunately, the sky is not a reliable indicator. We have to look at the scene.

Is the sun shining on the scene?

Sun shining

Sun hiding

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This is really nice, but unfortunately it doesn't capture the whole story: if we never see the sun, how can we know whether it's shining directly on the scene, or hiding behind a cloud? This is really important, because the scene will be lit very differently whether the sun is shining directly on it, or not. Unfortunately, the sky is not a reliable indicator. We have to look at the scene.

Mean scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We use a set of features which rely on the mean image computed over all the images in the sequence. As you can see, it looks a lot like an image where the sun is not shining on the scene.

Mean scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We use a set of features which rely on the mean image computed over all the images in the sequence. As you can see, it looks a lot like an image where the sun is not shining on the scene.

Mean scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We use a set of features which rely on the mean image computed over all the images in the sequence. As you can see, it looks a lot like an image where the sun is not shining on the scene.

Is the sun shining on the scene?

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

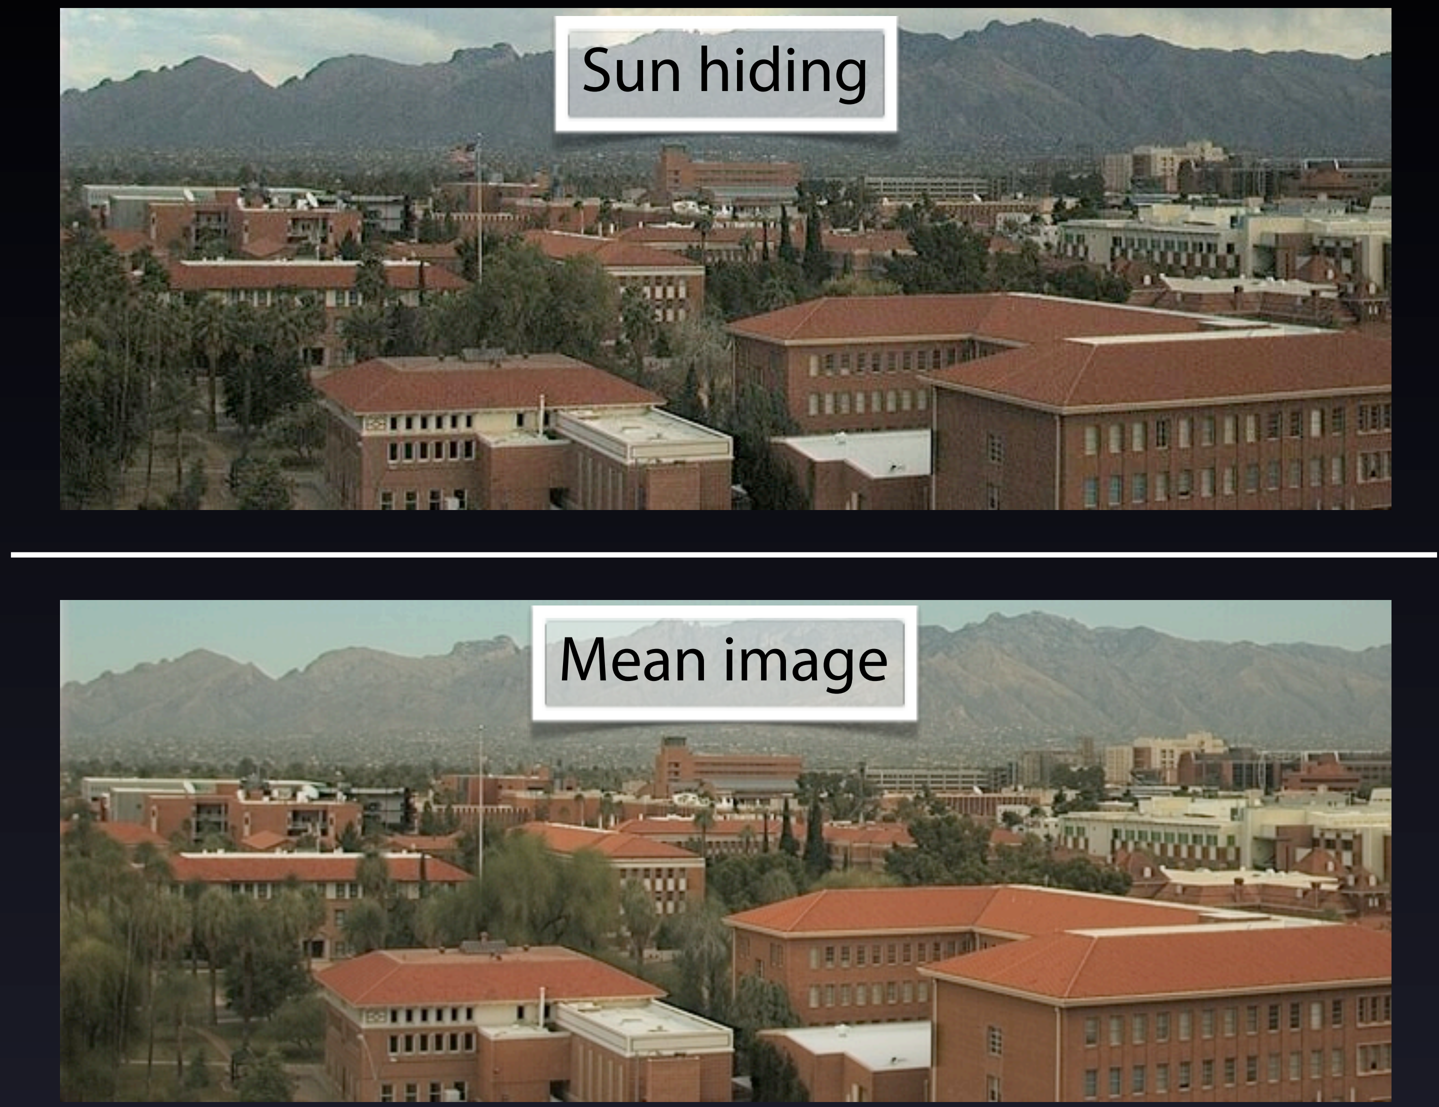

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

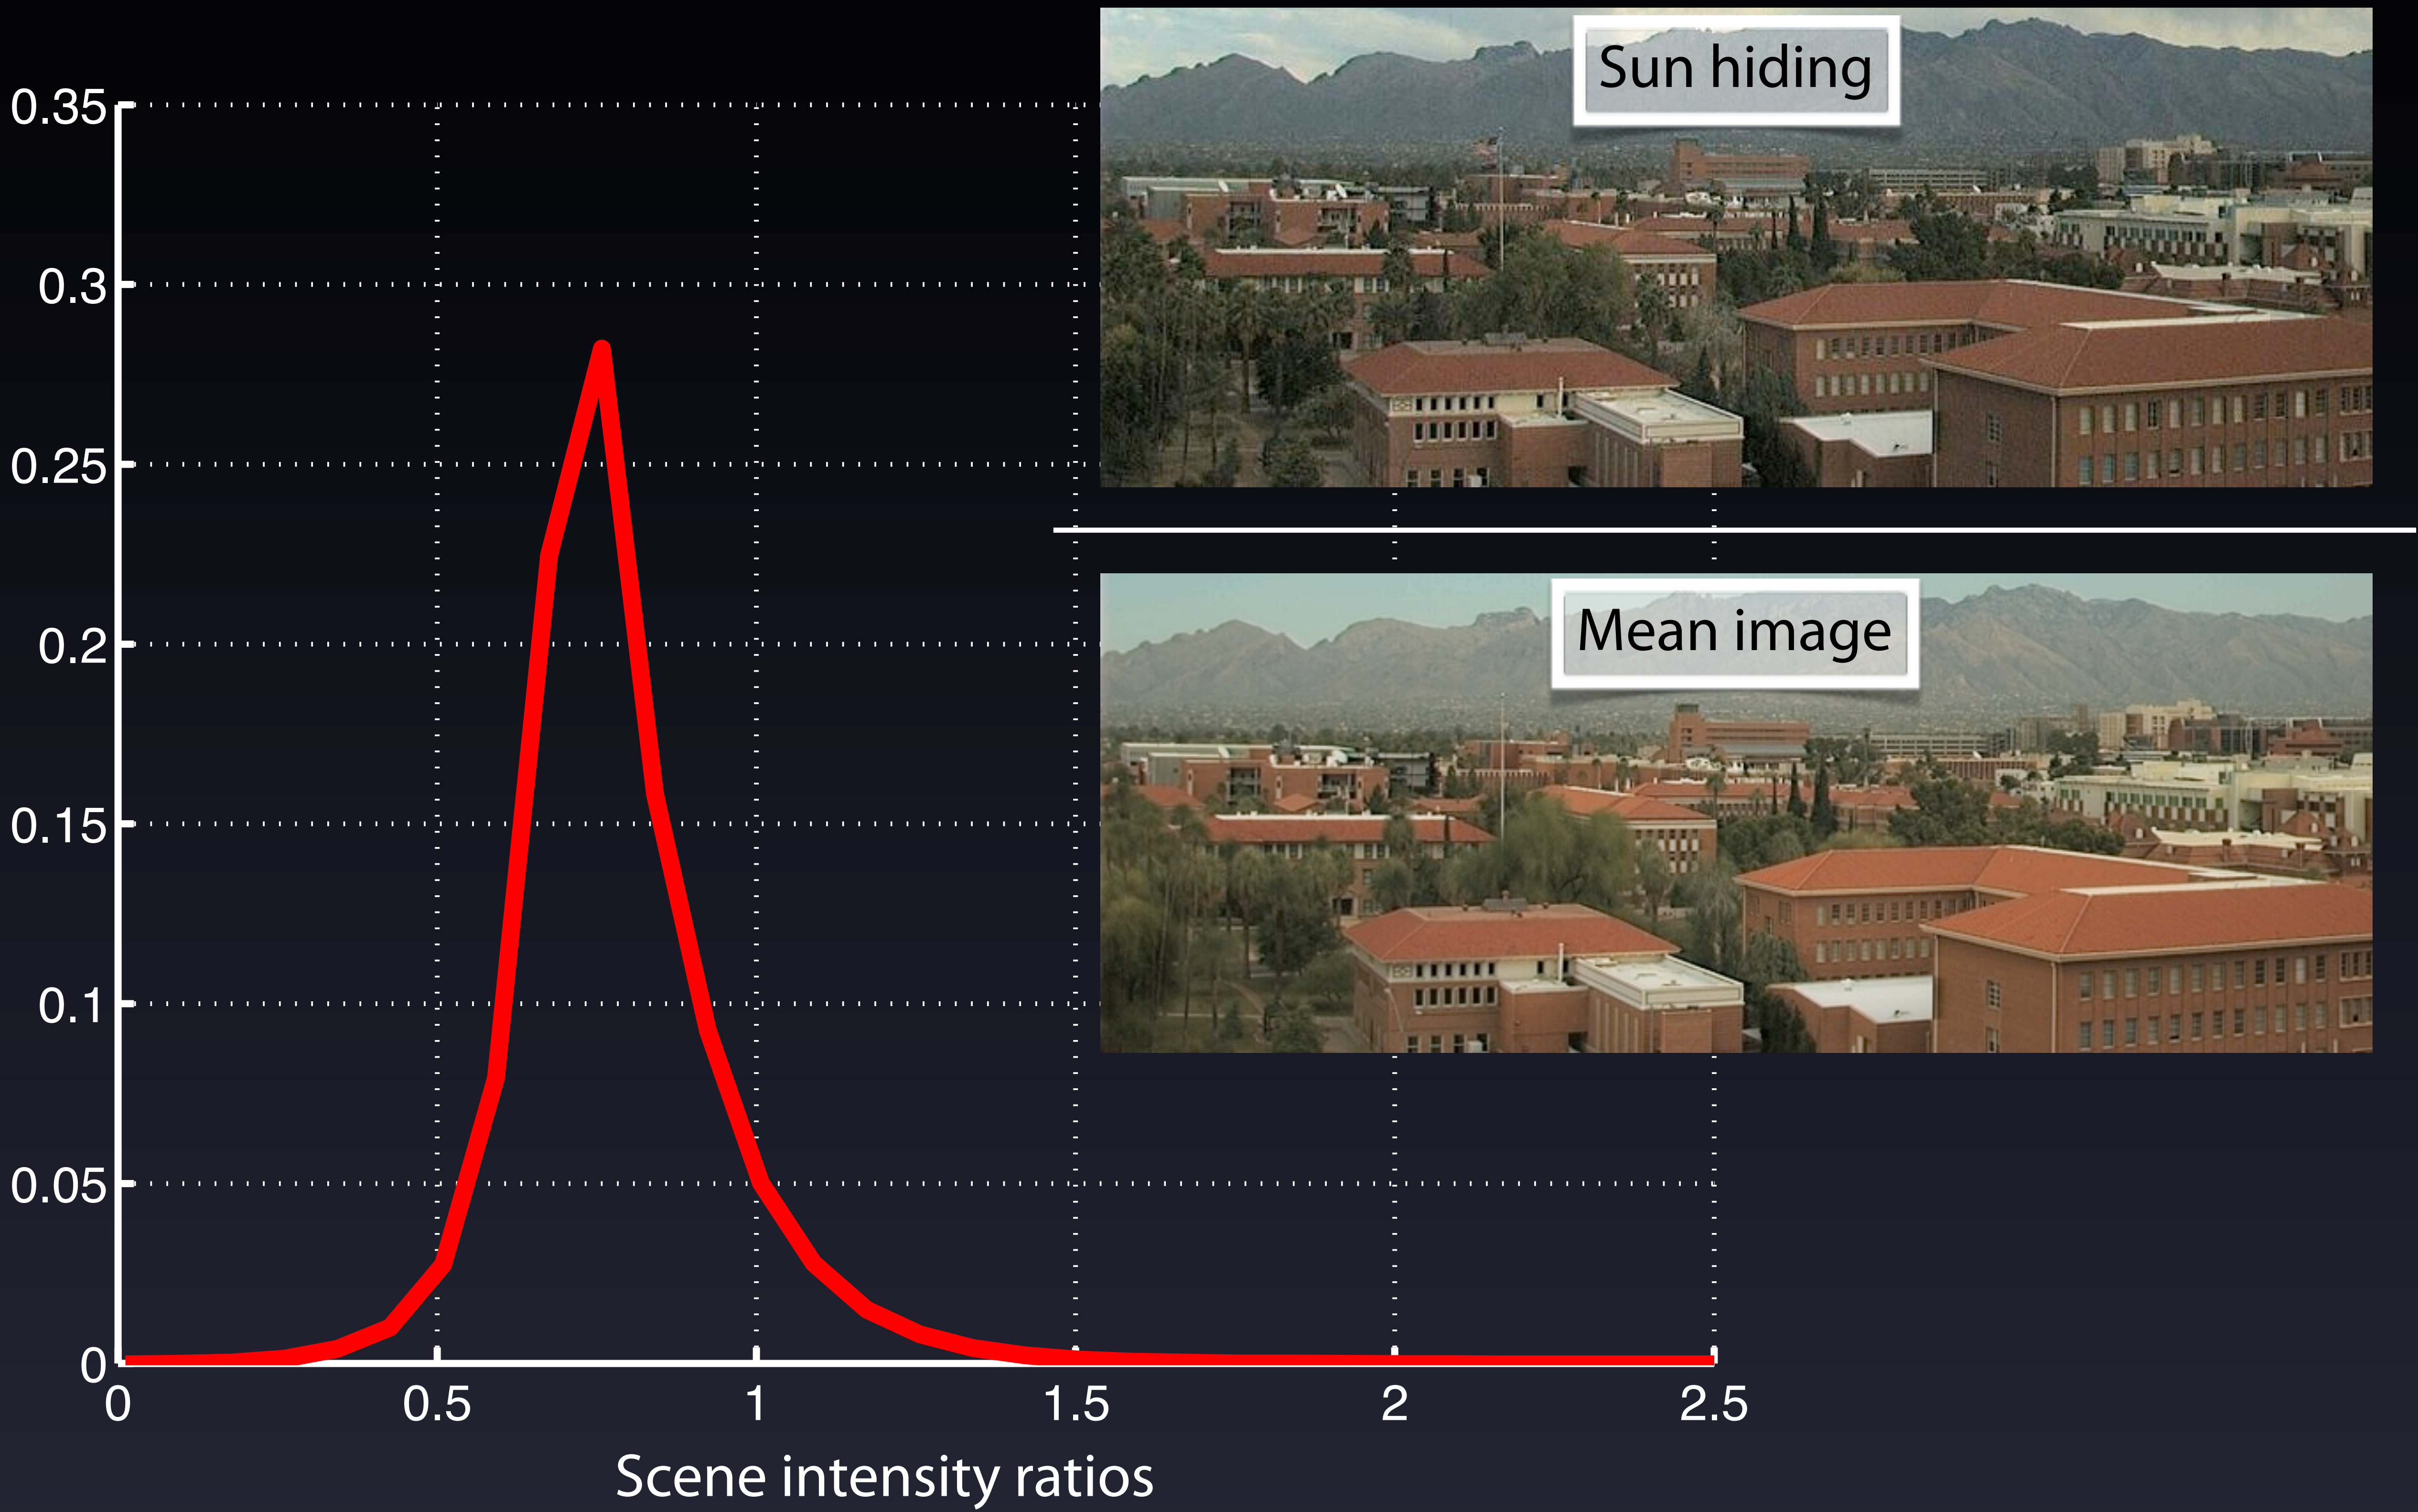

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let’s take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

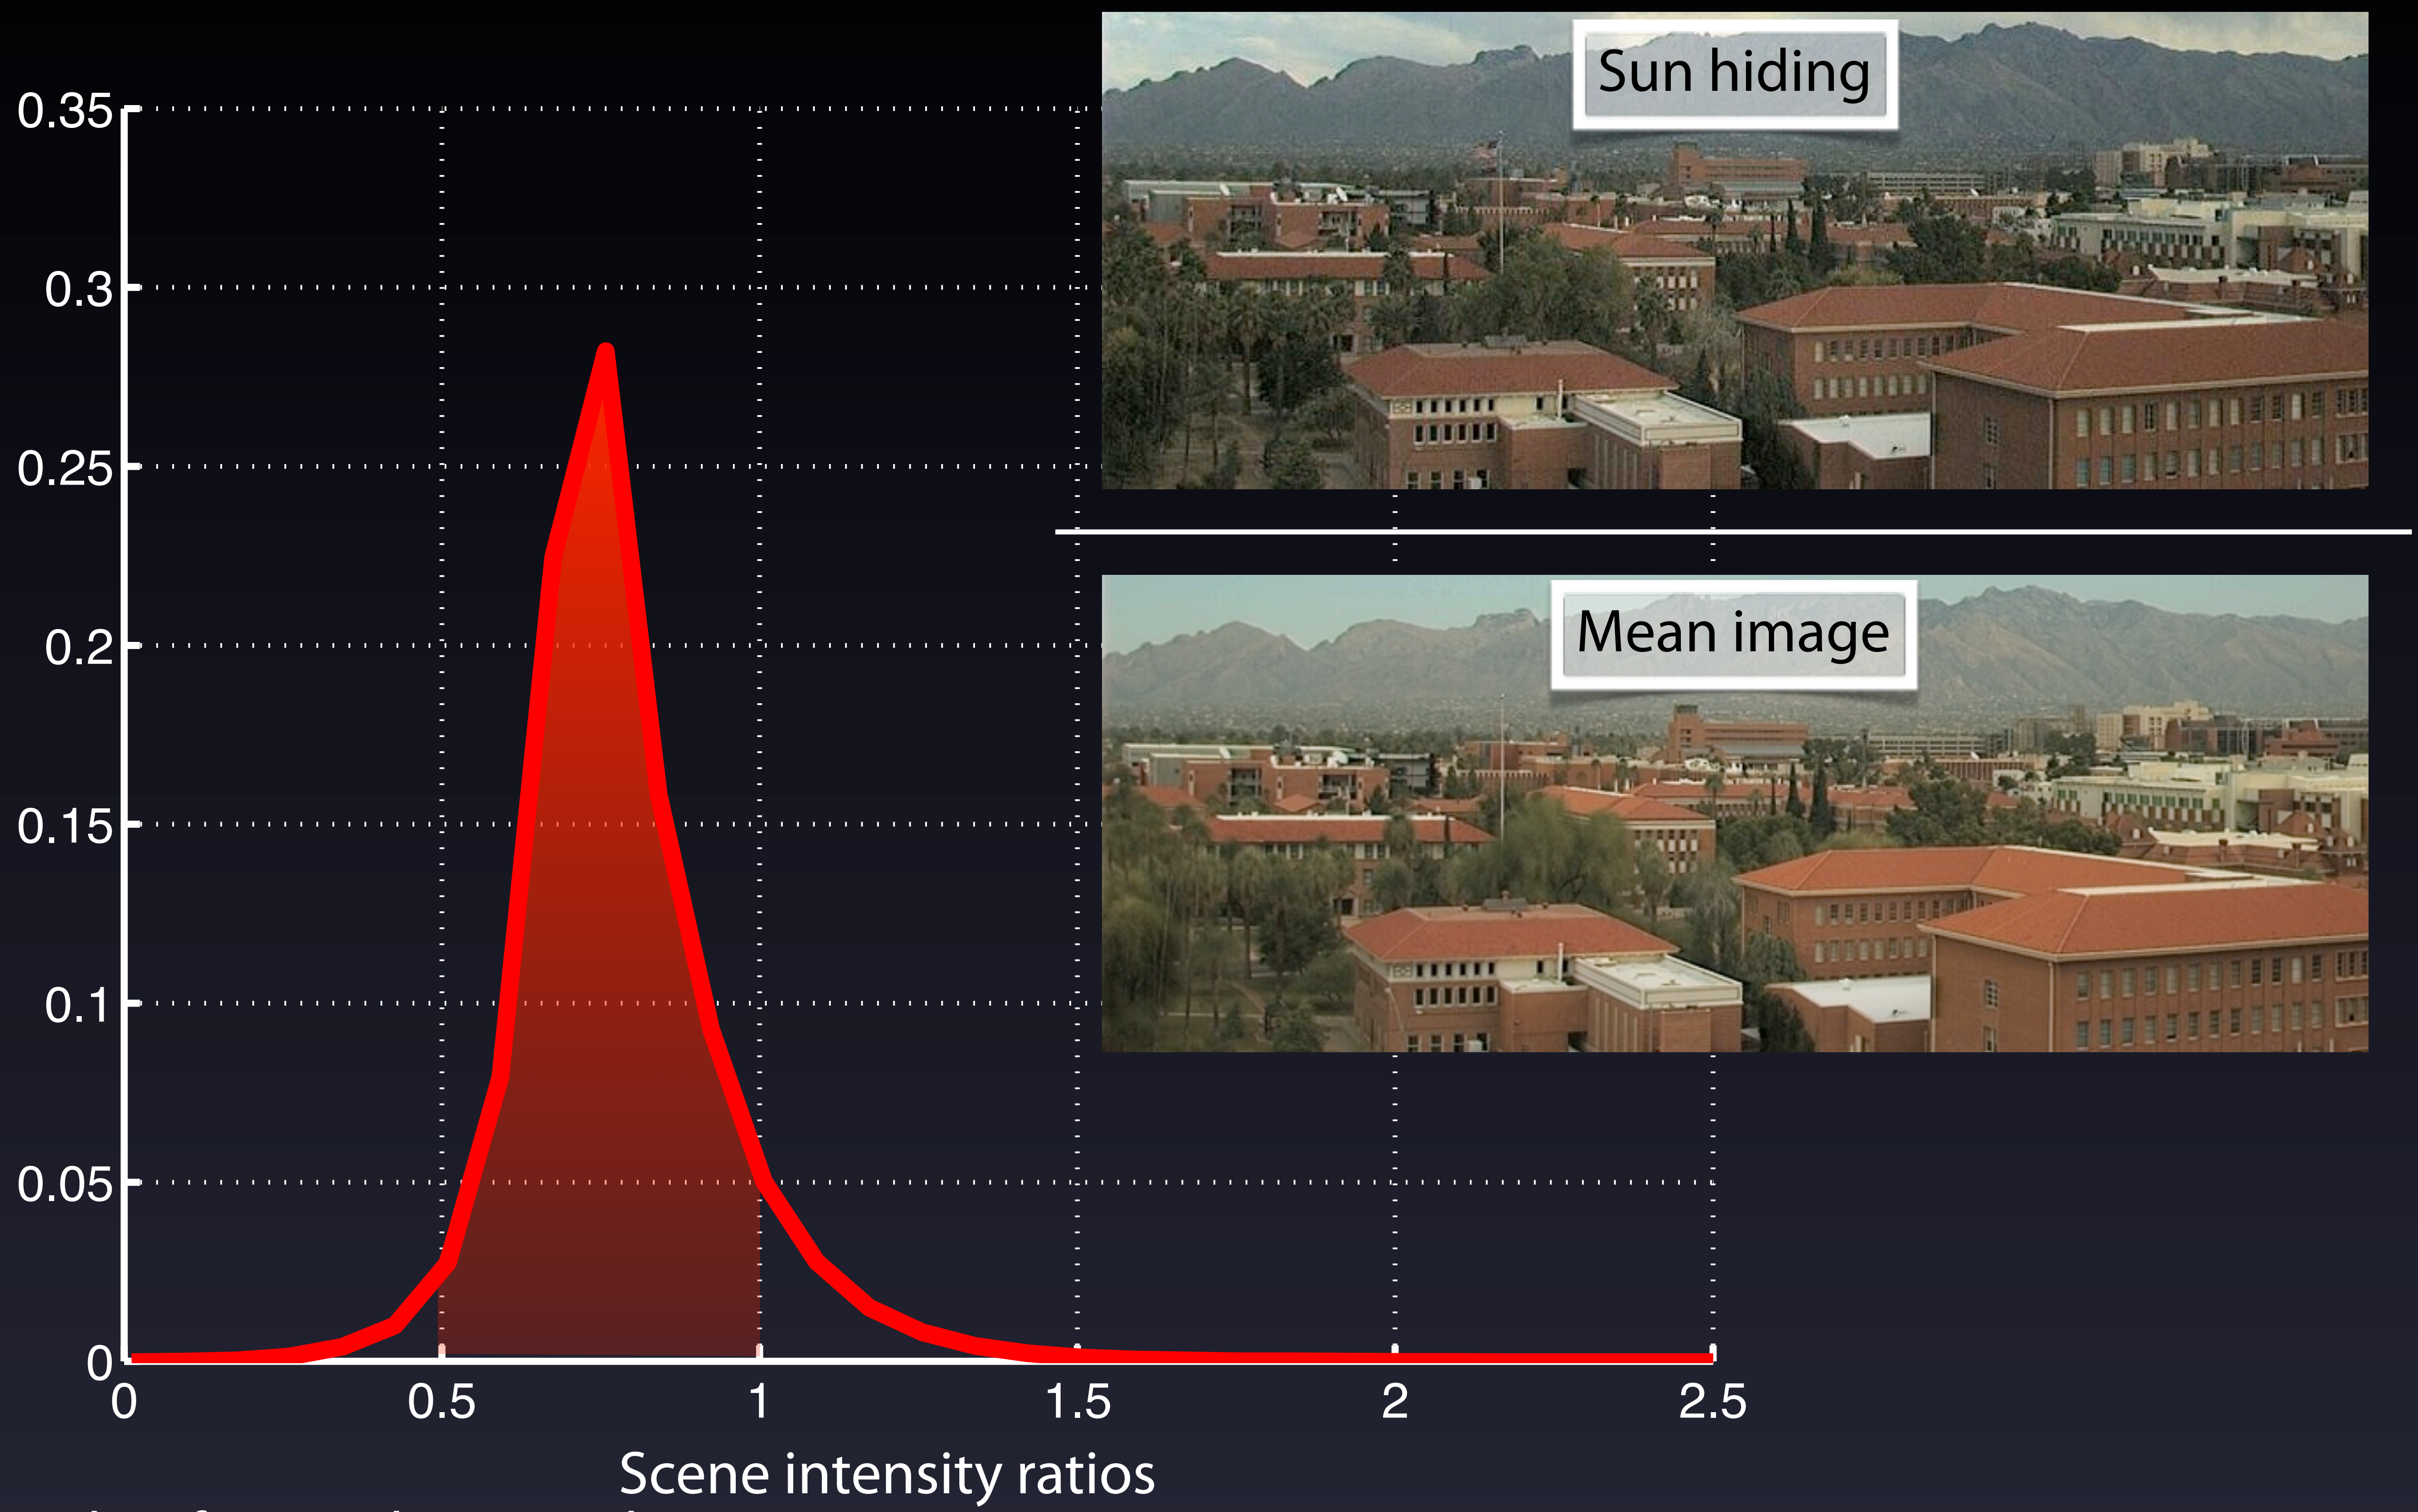

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.

We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

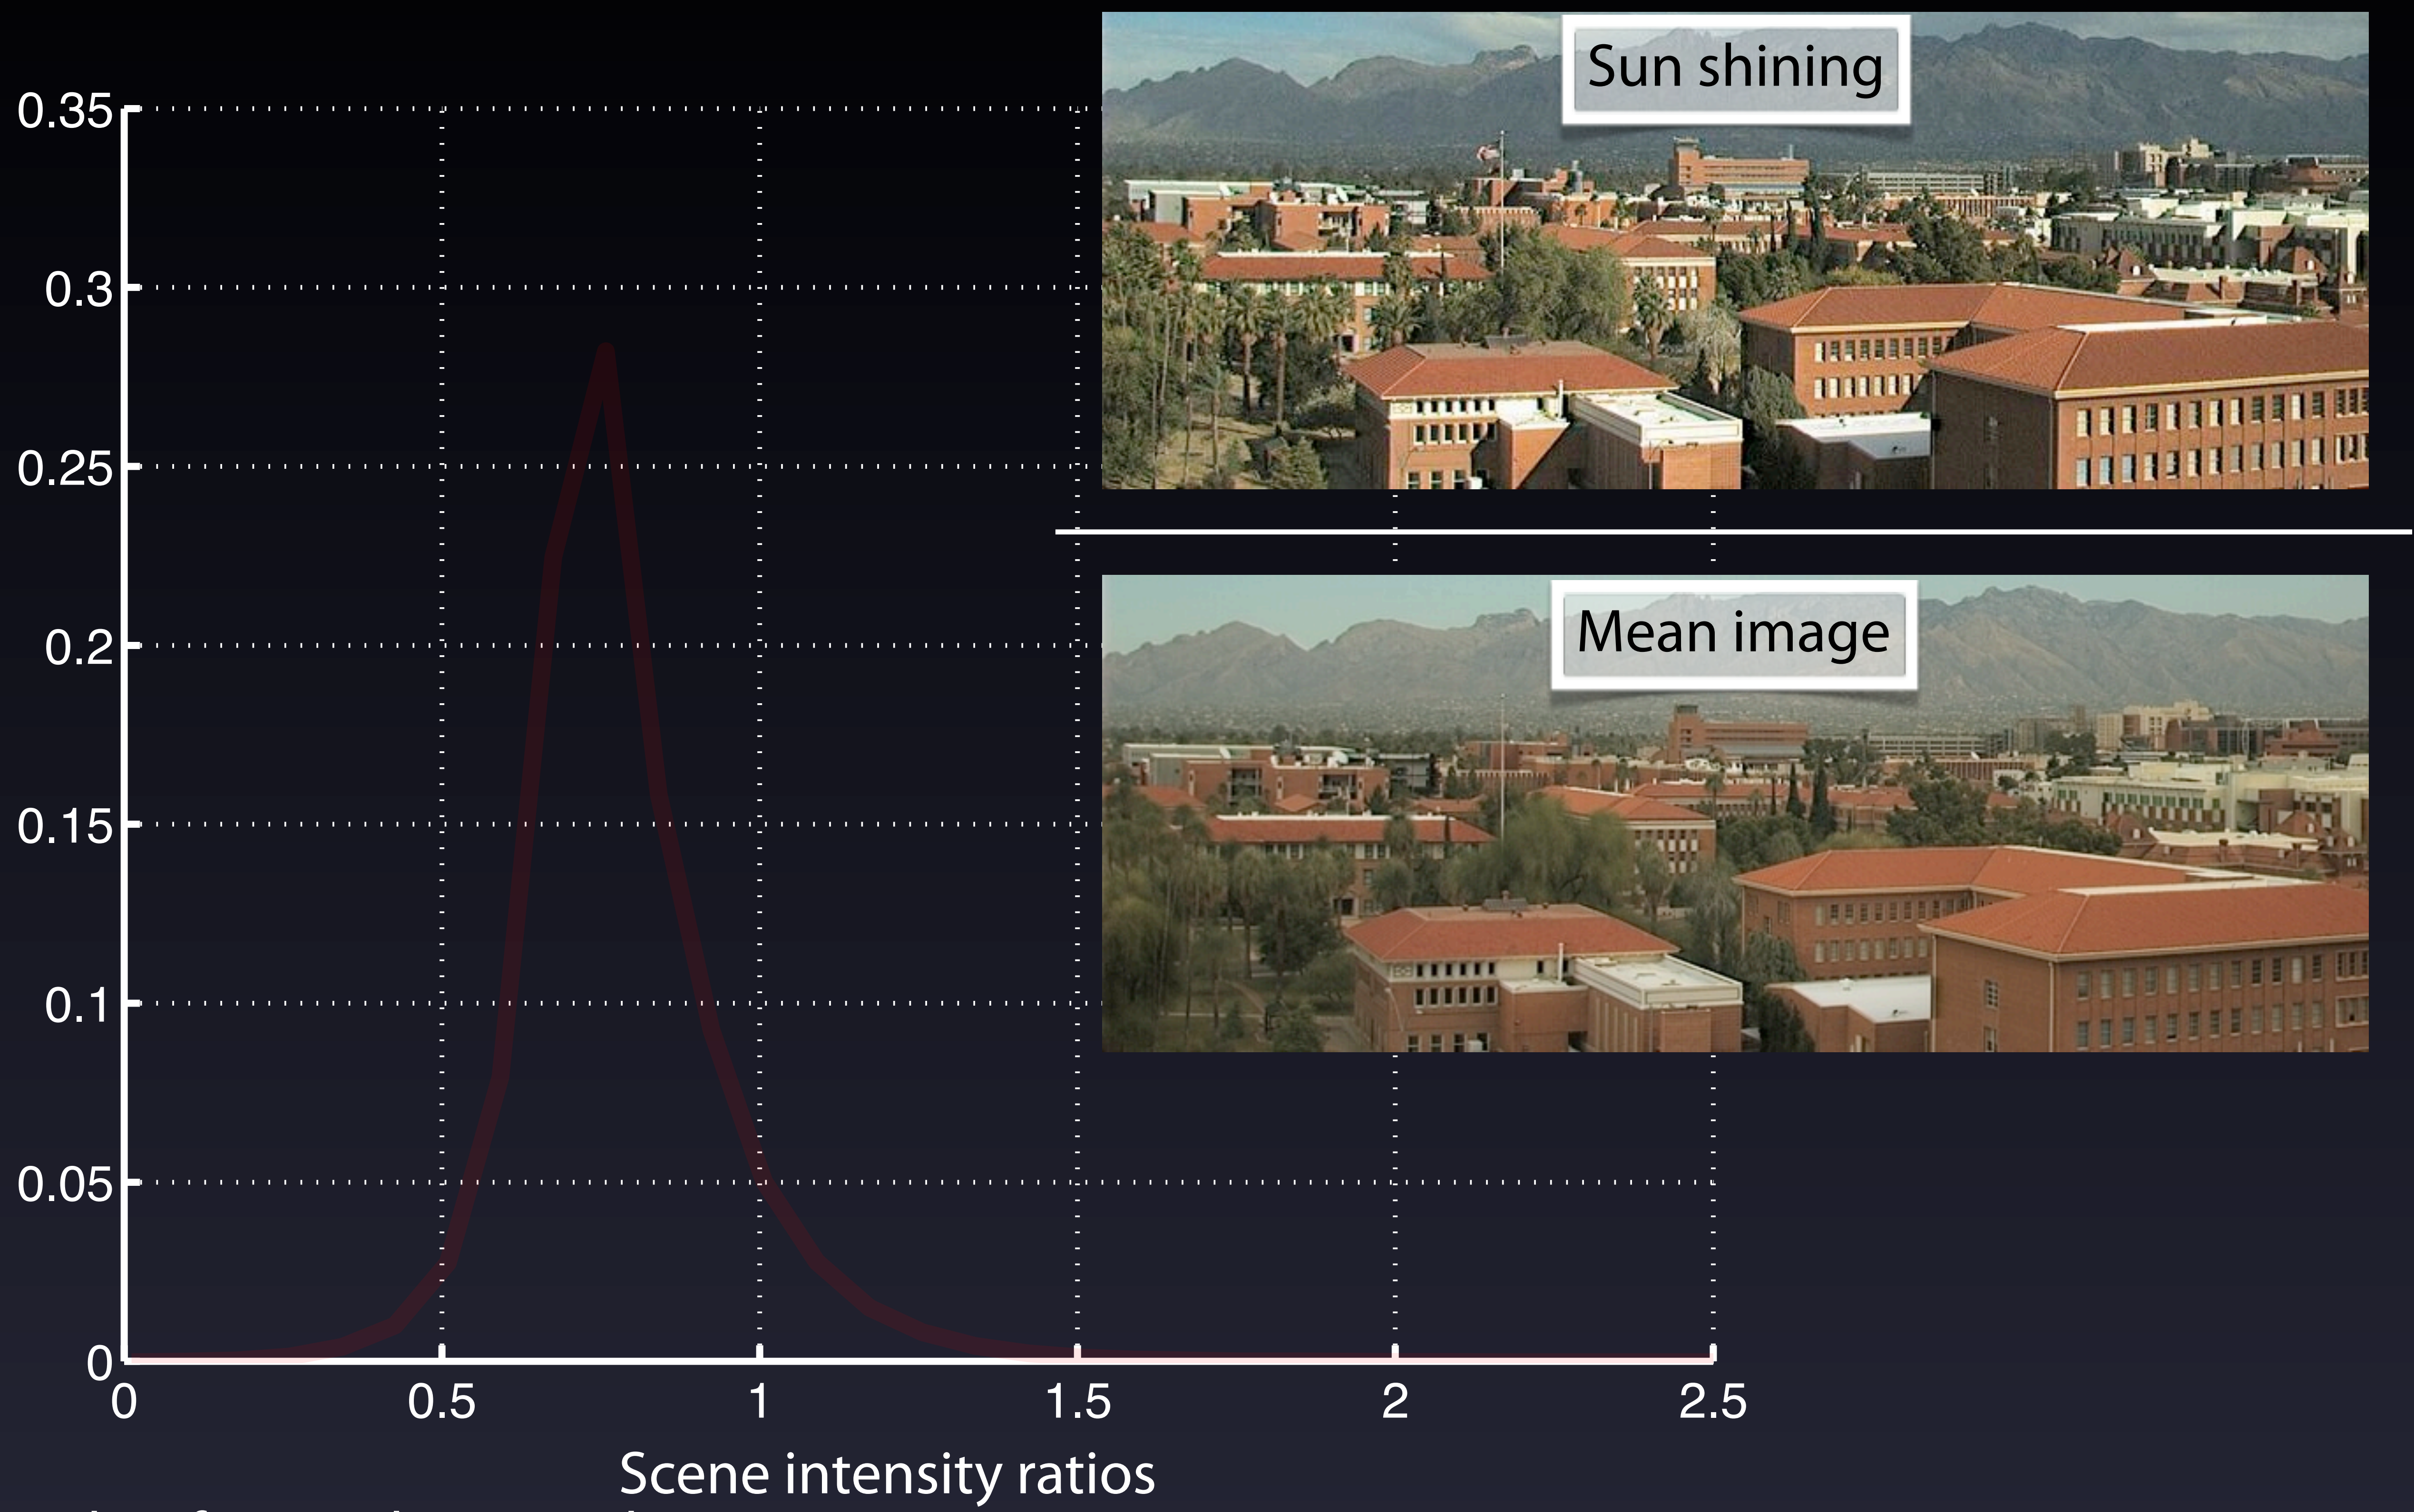

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.

We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

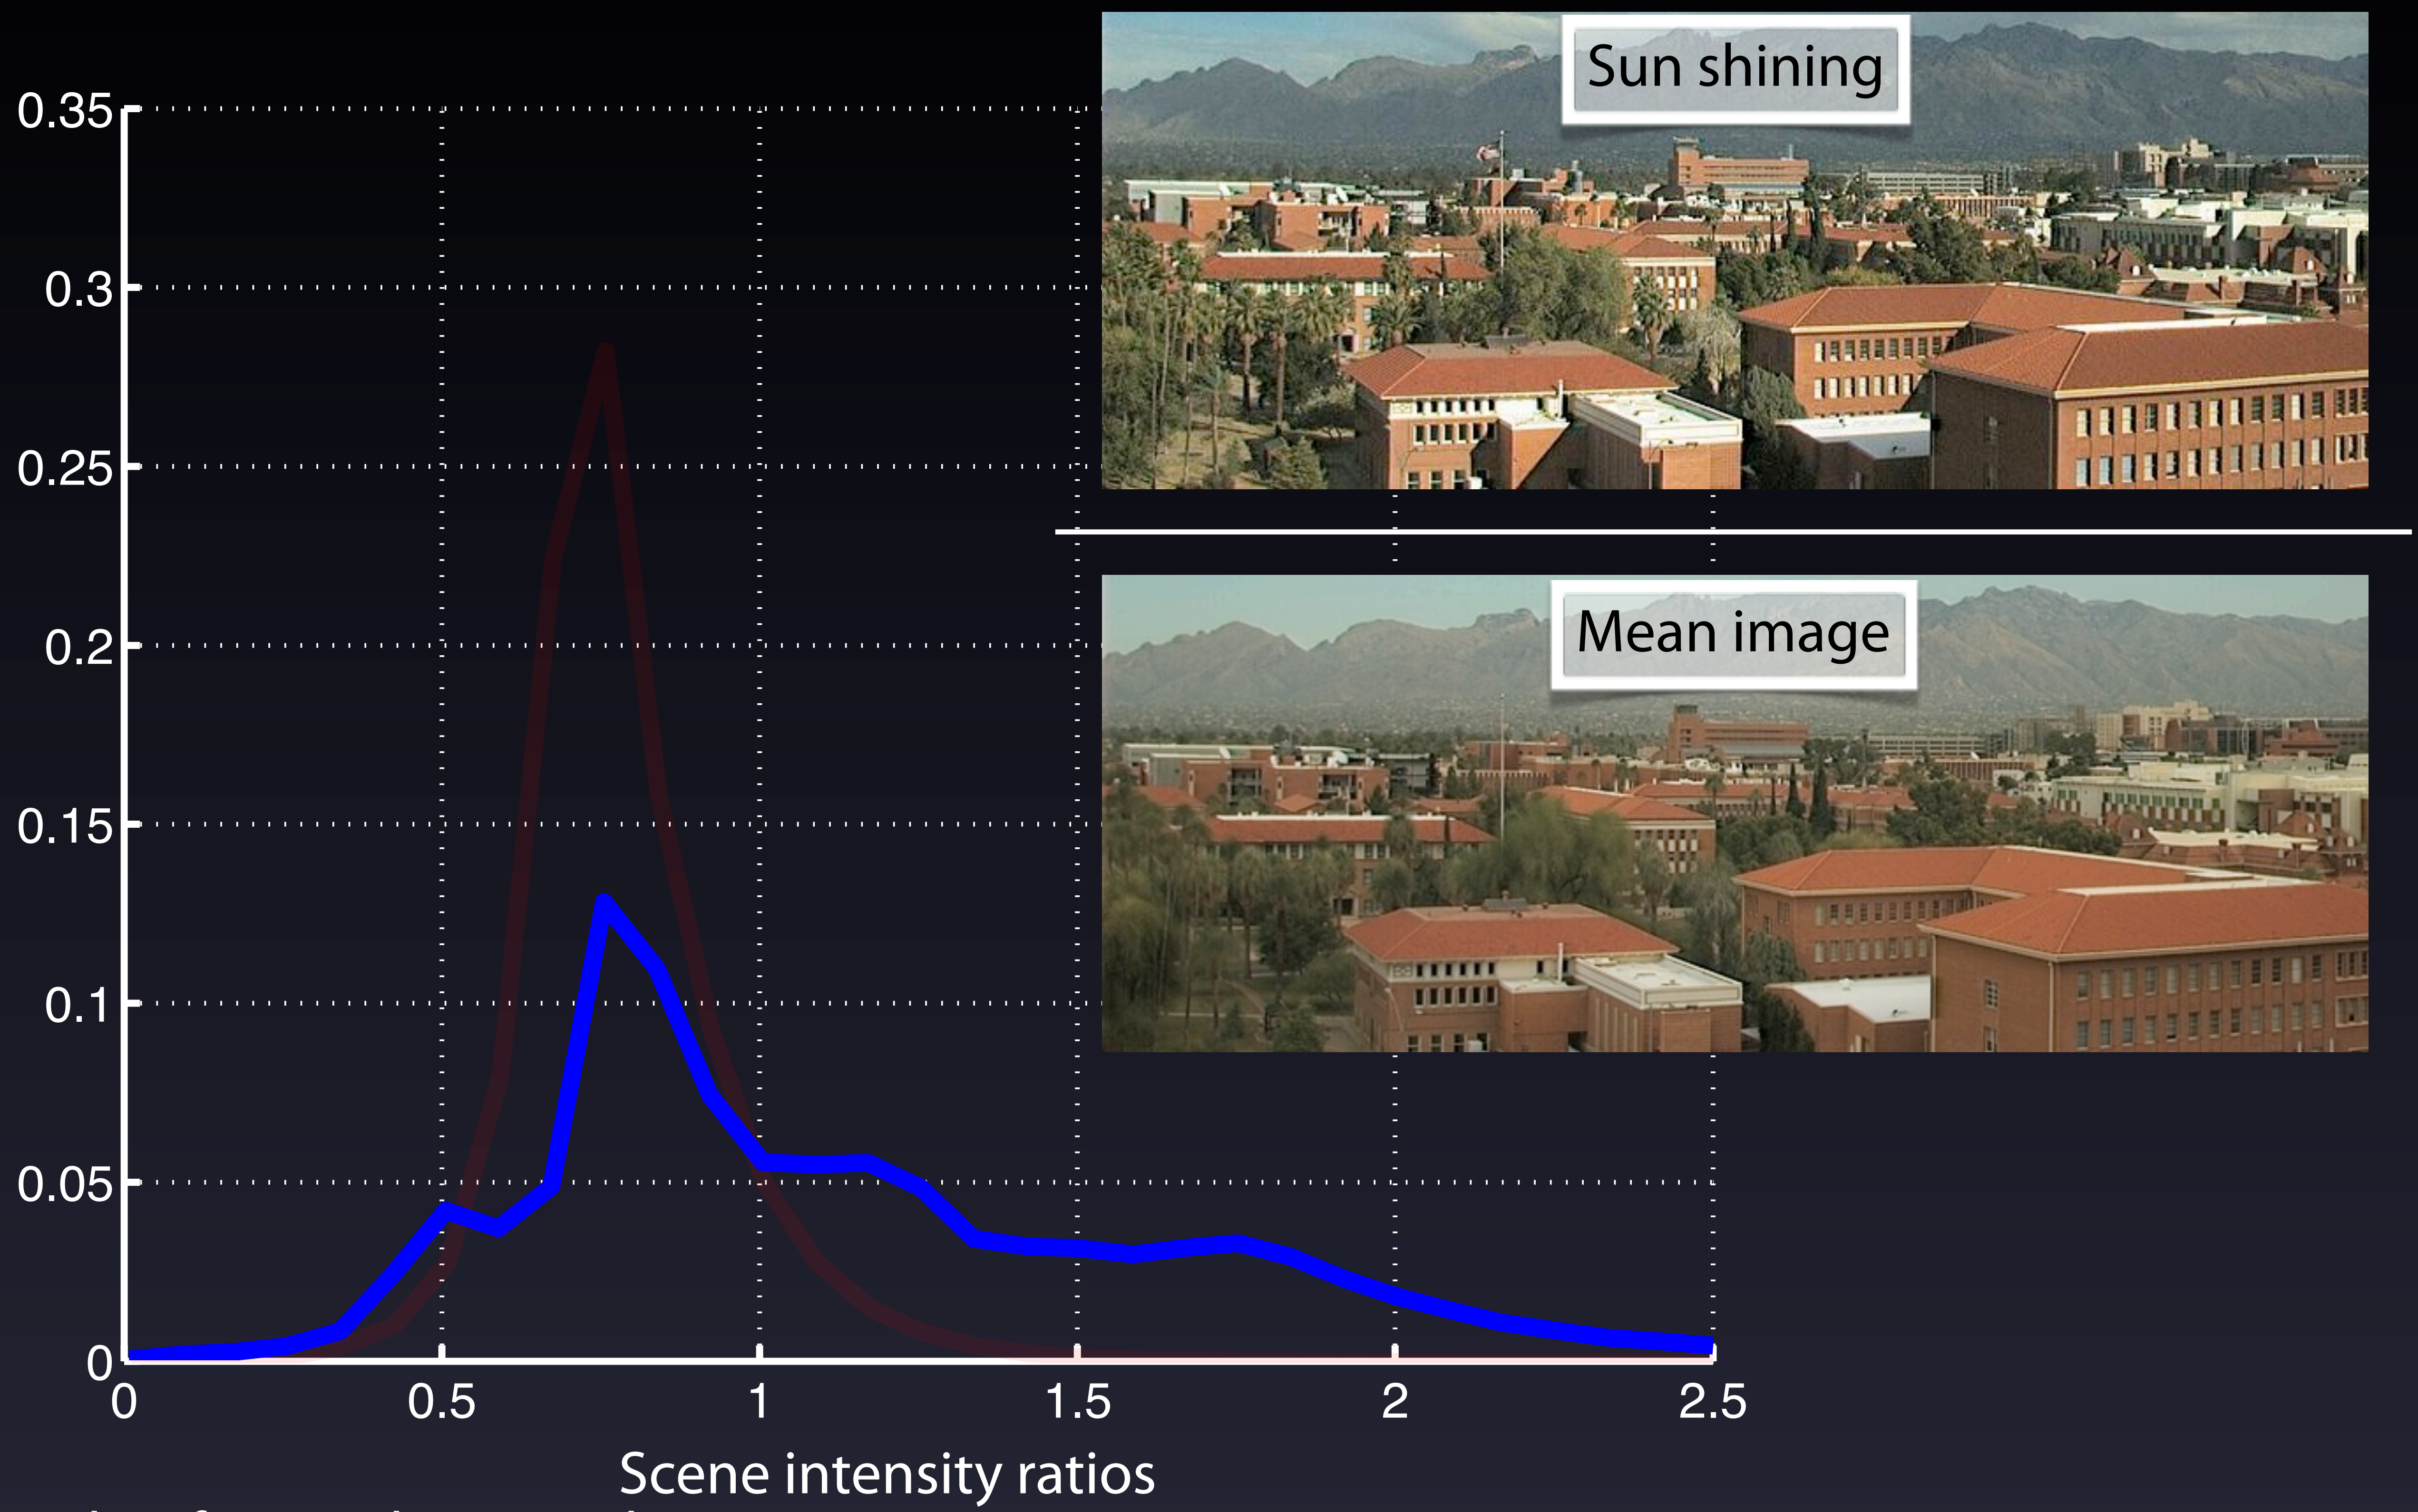

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

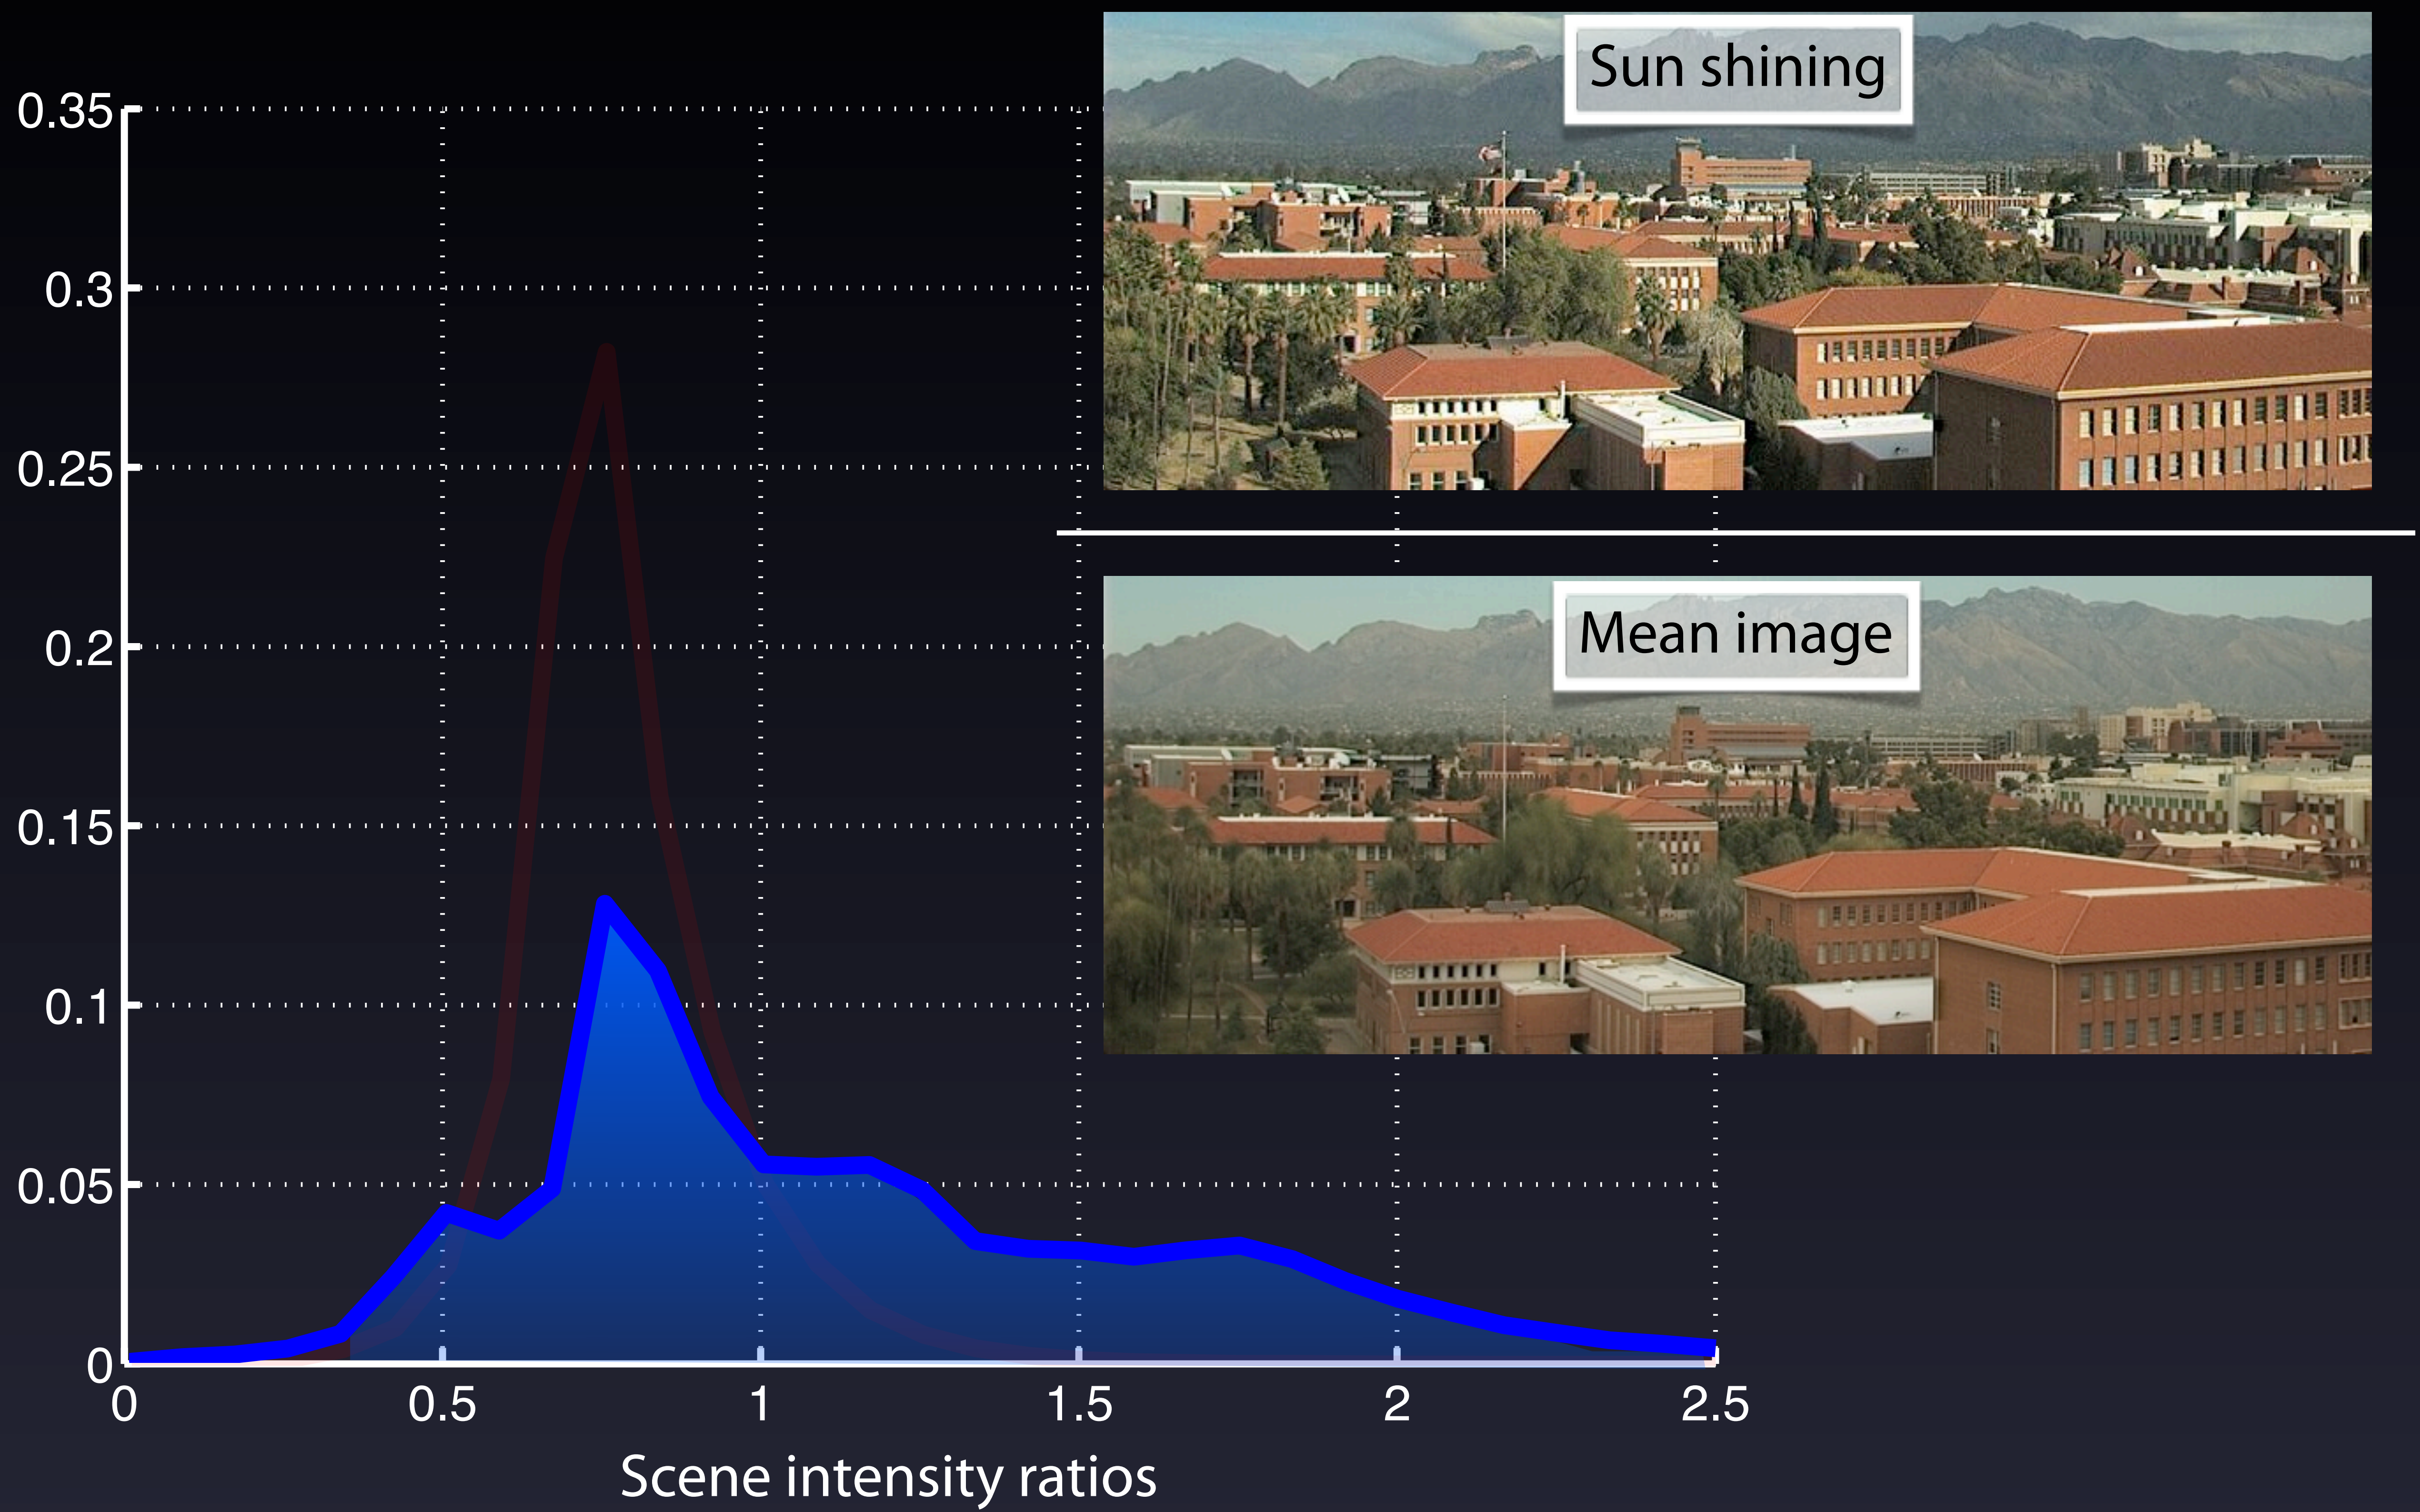

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.

We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Is the sun shining on the scene?

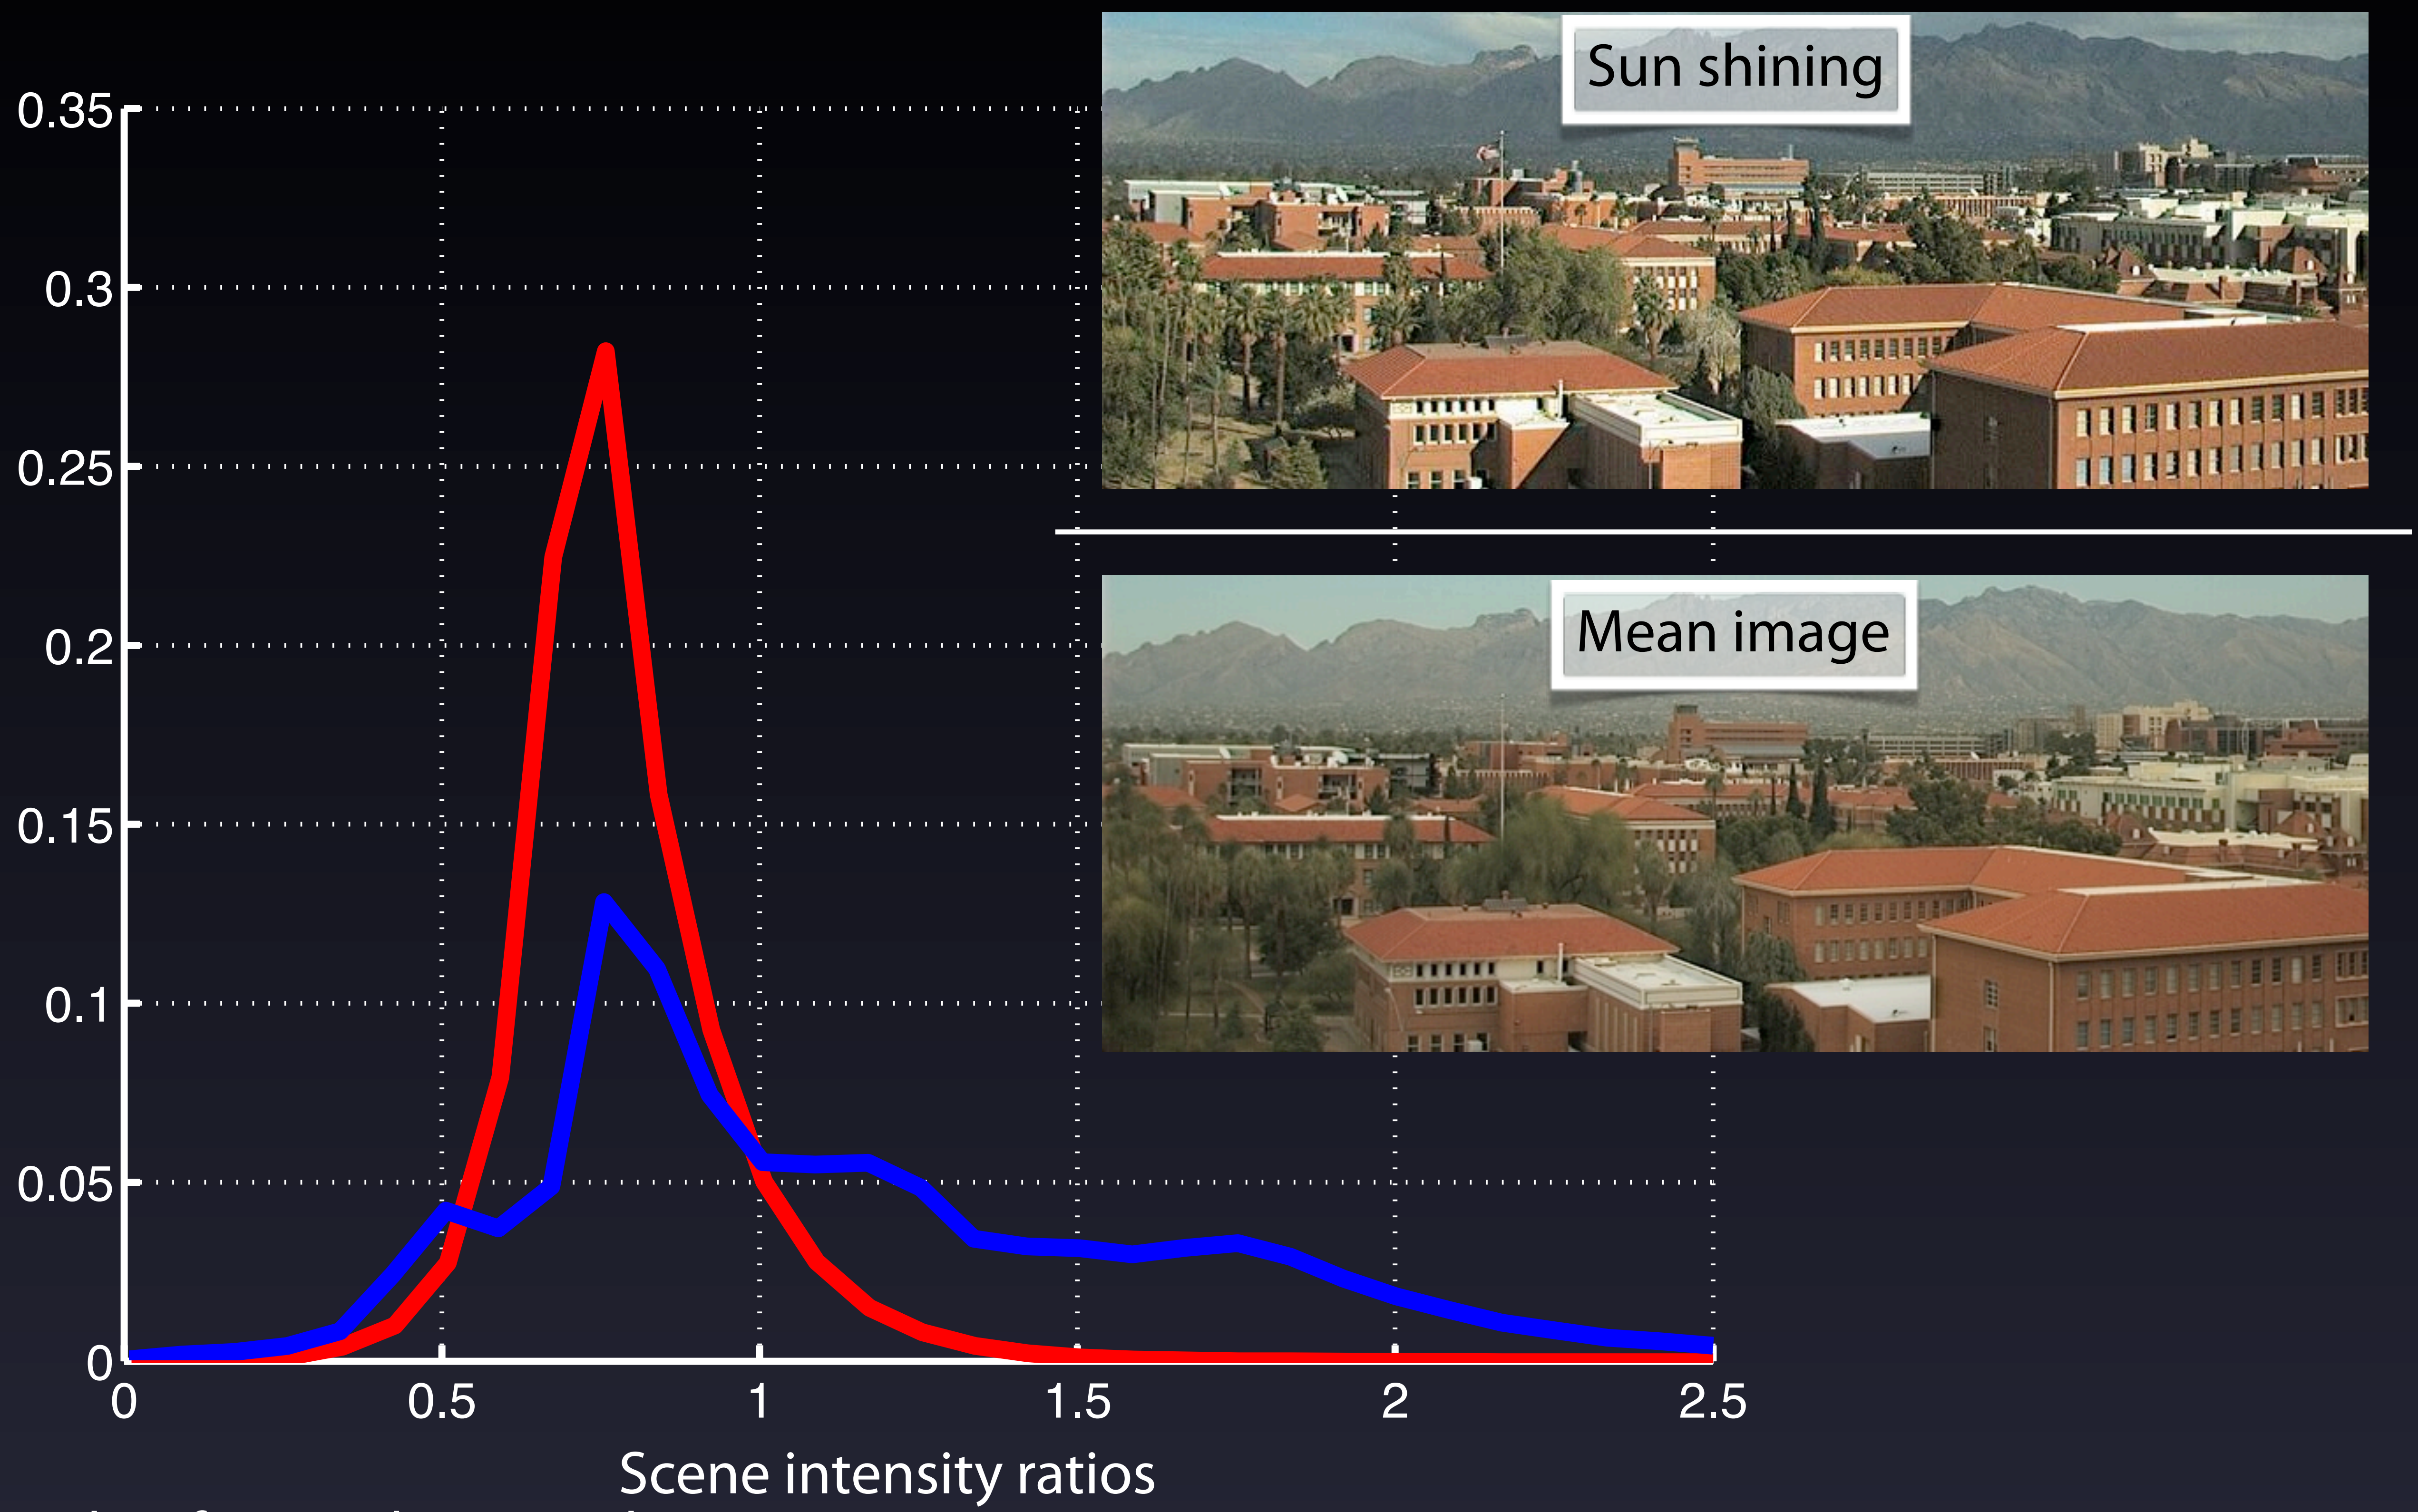

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Based on that observation, let's take the image where the sun is not shining on the scene, divide it by the mean image, and compute the histogram of these ratios, we get a distribution that has a very narrow peak. If we do the same for the image where the sun IS shining on the scene, we get a very different curve, with a much broader range.
We use this feature in a nearest-neighbor classifier trained on labelled images to determine whether the sun is hidden behind a cloud or not.

Illumination from webcams

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Now that we've estimated the illumination conditions from a sequence, namely the sun position and whether it's shining on the scene or not, sky turbidity and cloud cover, how can we use this to generate our synthetic environment map? Our goal is to find the color and intensity of the light coming in from all directions, even if the camera really only sees a very small portion of it.

Illumination from webcams

- sun position

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Now that we’ve estimated the illumination conditions from a sequence, namely the sun position and whether it’s shining on the scene or not, sky turbidity and cloud cover, how can we use this to generate our synthetic environment map? Our goal is to find the color and intensity of the light coming in from all directions, even if the camera really only sees a very small portion of it.

Illumination from webcams

- sun position
- sun shining or not

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Now that we've estimated the illumination conditions from a sequence, namely the sun position and whether it's shining on the scene or not, sky turbidity and cloud cover, how can we use this to generate our synthetic environment map? Our goal is to find the color and intensity of the light coming in from all directions, even if the camera really only sees a very small portion of it.

Illumination from webcams

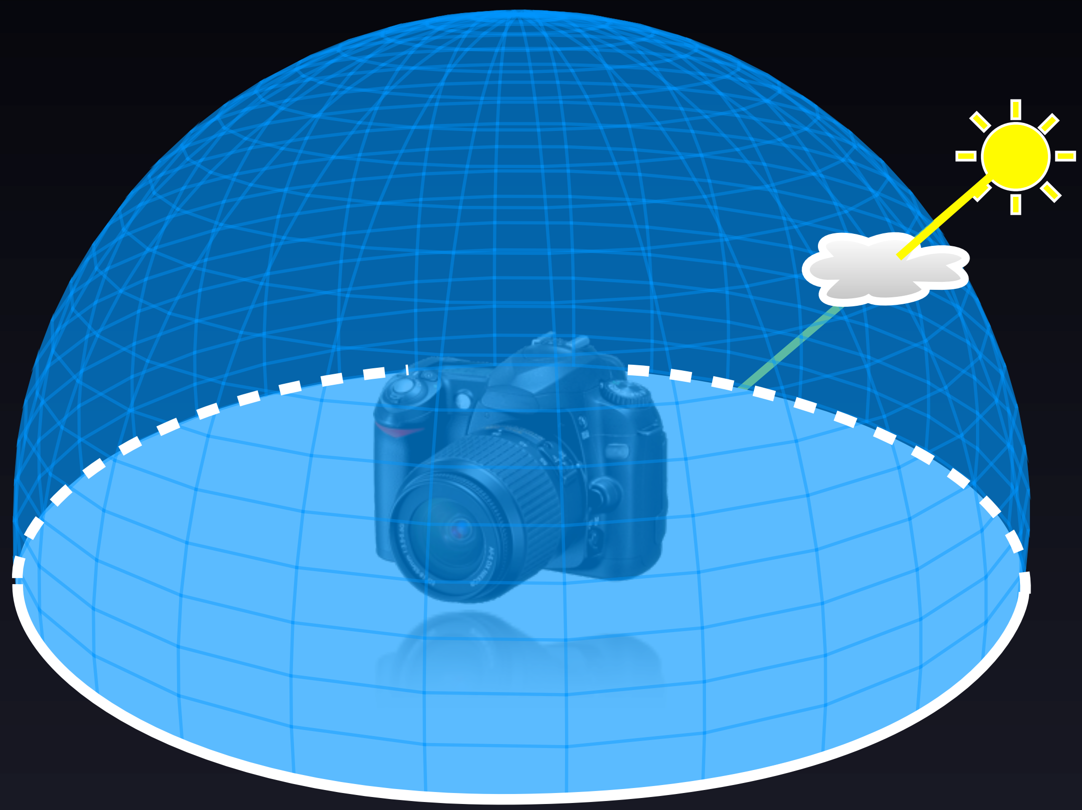

- sun position
- sun shining or not
- sky turbidity

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Now that we've estimated the illumination conditions from a sequence, namely the sun position and whether it's shining on the scene or not, sky turbidity and cloud cover, how can we use this to generate our synthetic environment map? Our goal is to find the color and intensity of the light coming in from all directions, even if the camera really only sees a very small portion of it.

Illumination from webcams

- sun position
- sun shining or not
- sky turbidity
- cloud segmentation

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Now that we've estimated the illumination conditions from a sequence, namely the sun position and whether it's shining on the scene or not, sky turbidity and cloud cover, how can we use this to generate our synthetic environment map? Our goal is to find the color and intensity of the light coming in from all directions, even if the camera really only sees a very small portion of it.

Sky

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

For the sky, we already know its turbidity, so all we need to do is to plug it in the Perez model, and we can recover its appearance everywhere. Notice that the sky model nicely captures the scattering due to the sun. Unfortunately, while we know the sun position, we don't know its intensity.

Sky

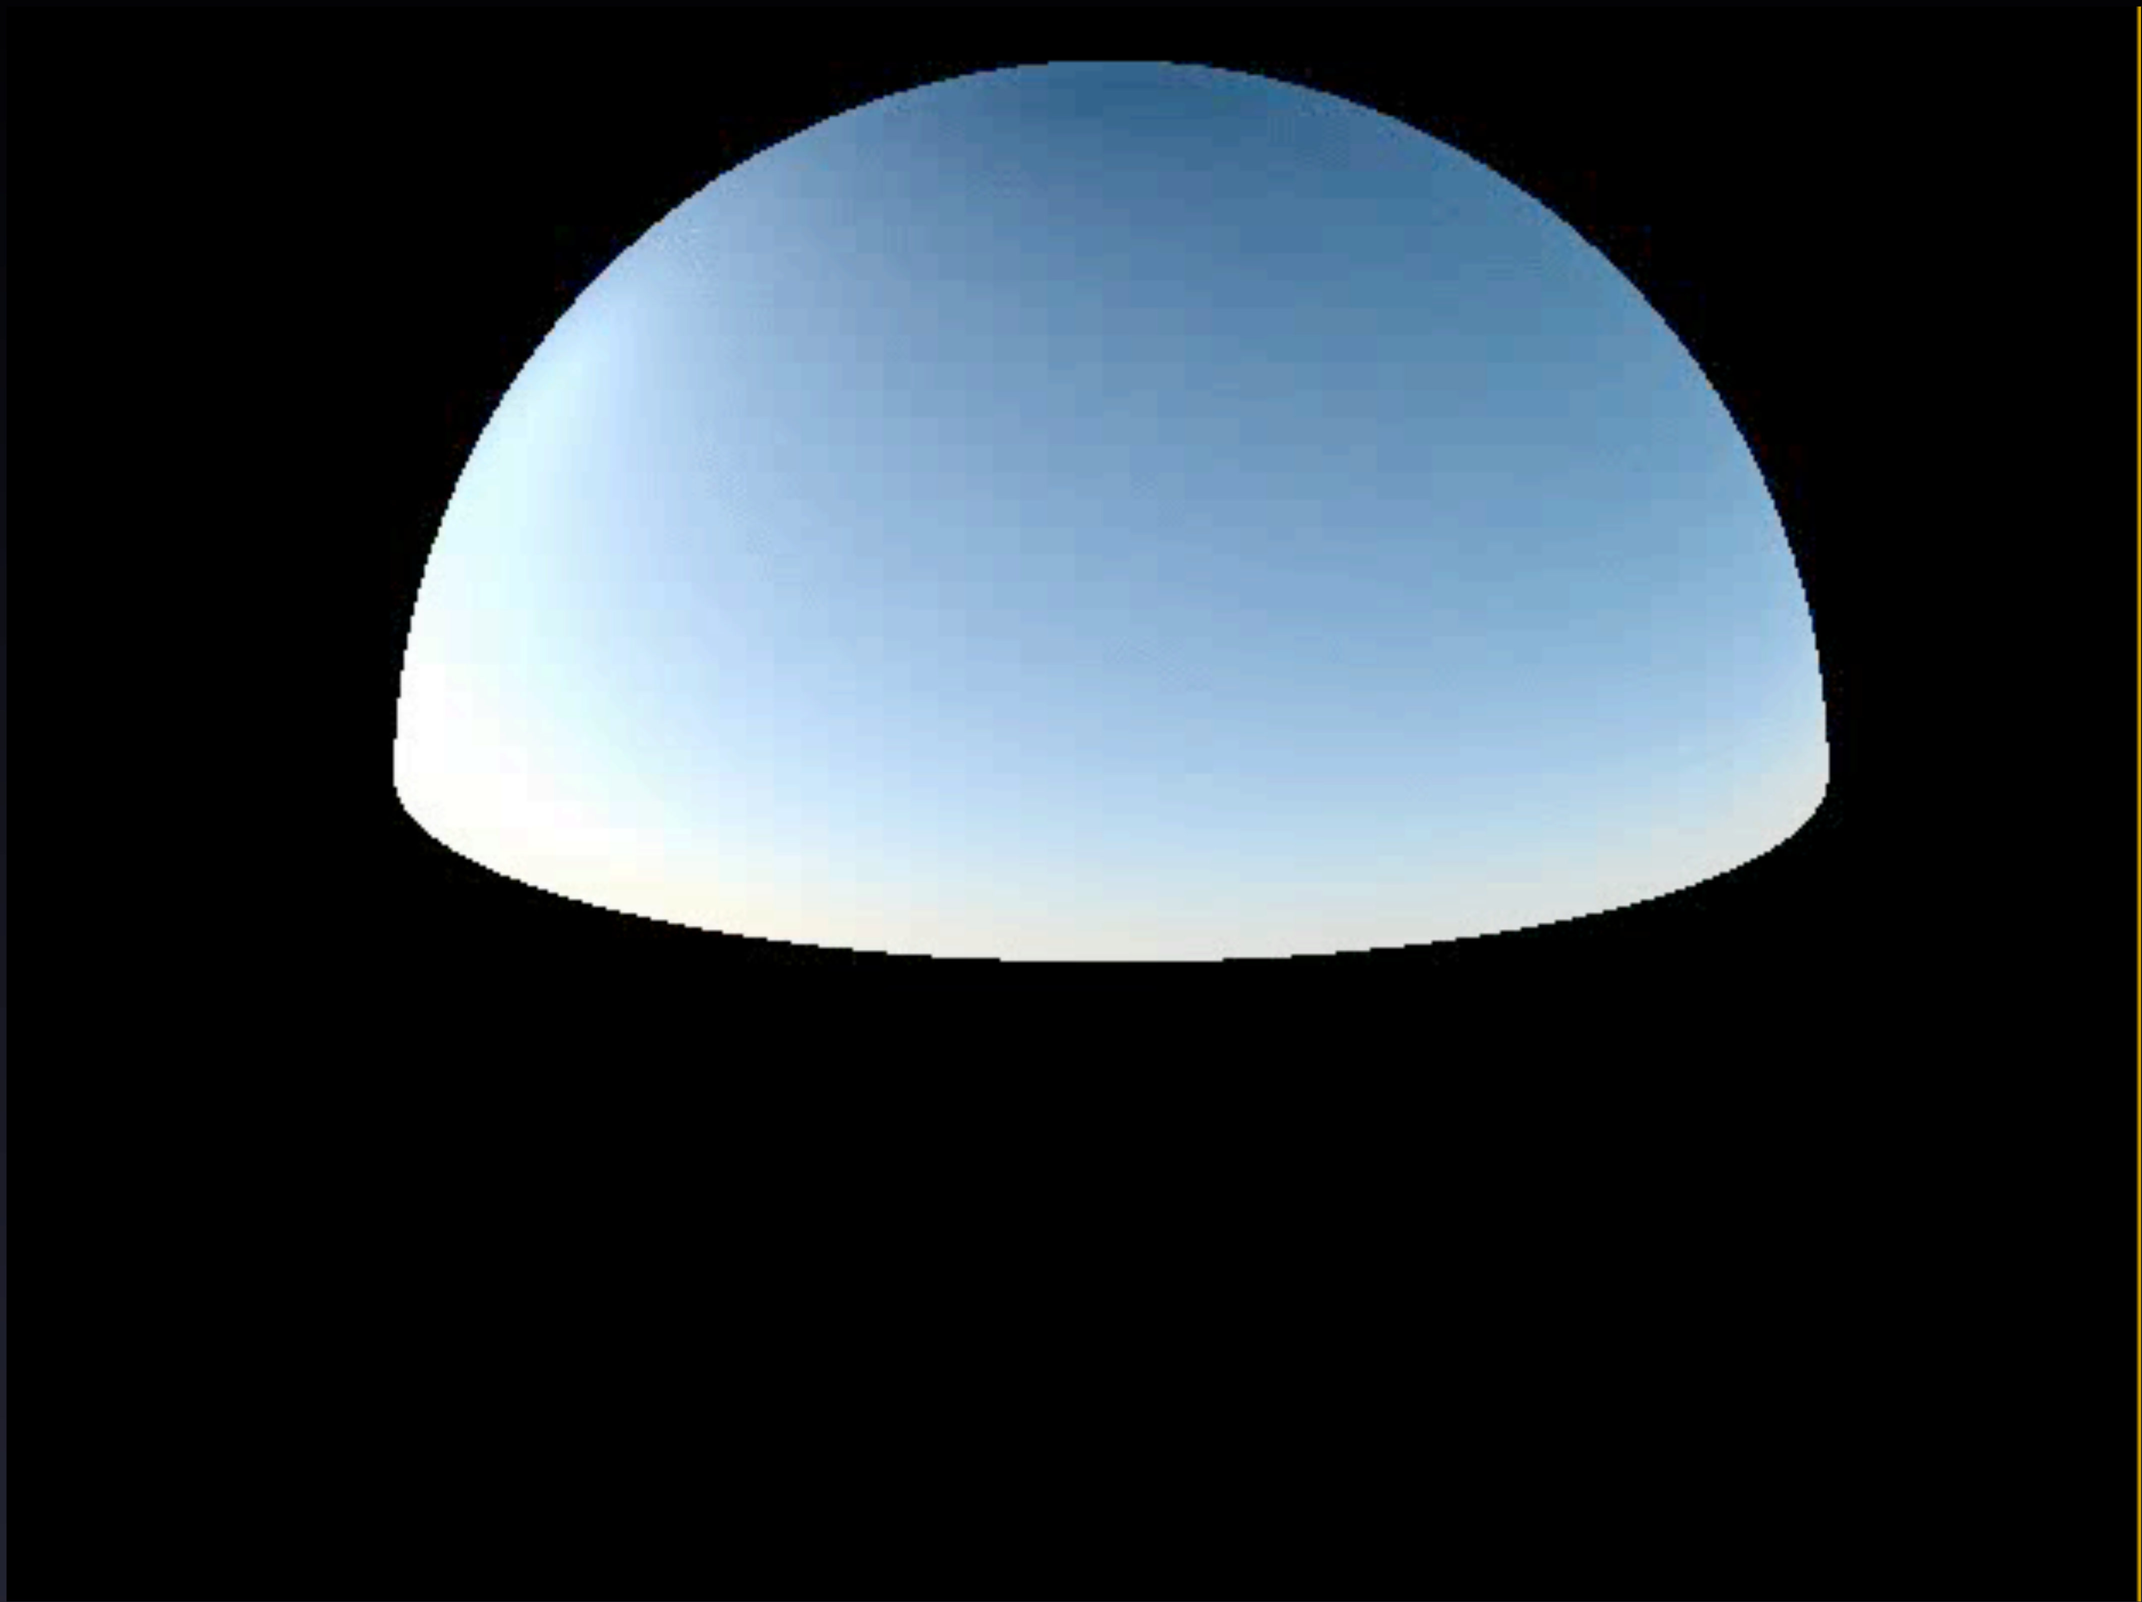

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

For the sky, we already know its turbidity, so all we need to do is to plug it in the Perez model, and we can recover its appearance everywhere. Notice that the sky model nicely captures the scattering due to the sun. Unfortunately, while we know the sun position, we don’t know its intensity.

Environment map – sun

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Stumpfel et al., 2004]

So we adopt a data-driven approach, and use the sky database of Debevec et al, which contains HDR pictures of the sky hemisphere taken over an entire day. We detect the sun at each frame where it is visible, and extract a small window around it. Note that I'm showing the sun positions in image space for visualization, but in practice we perform the computations in angular space to avoid any distortions.

Environment map – sun

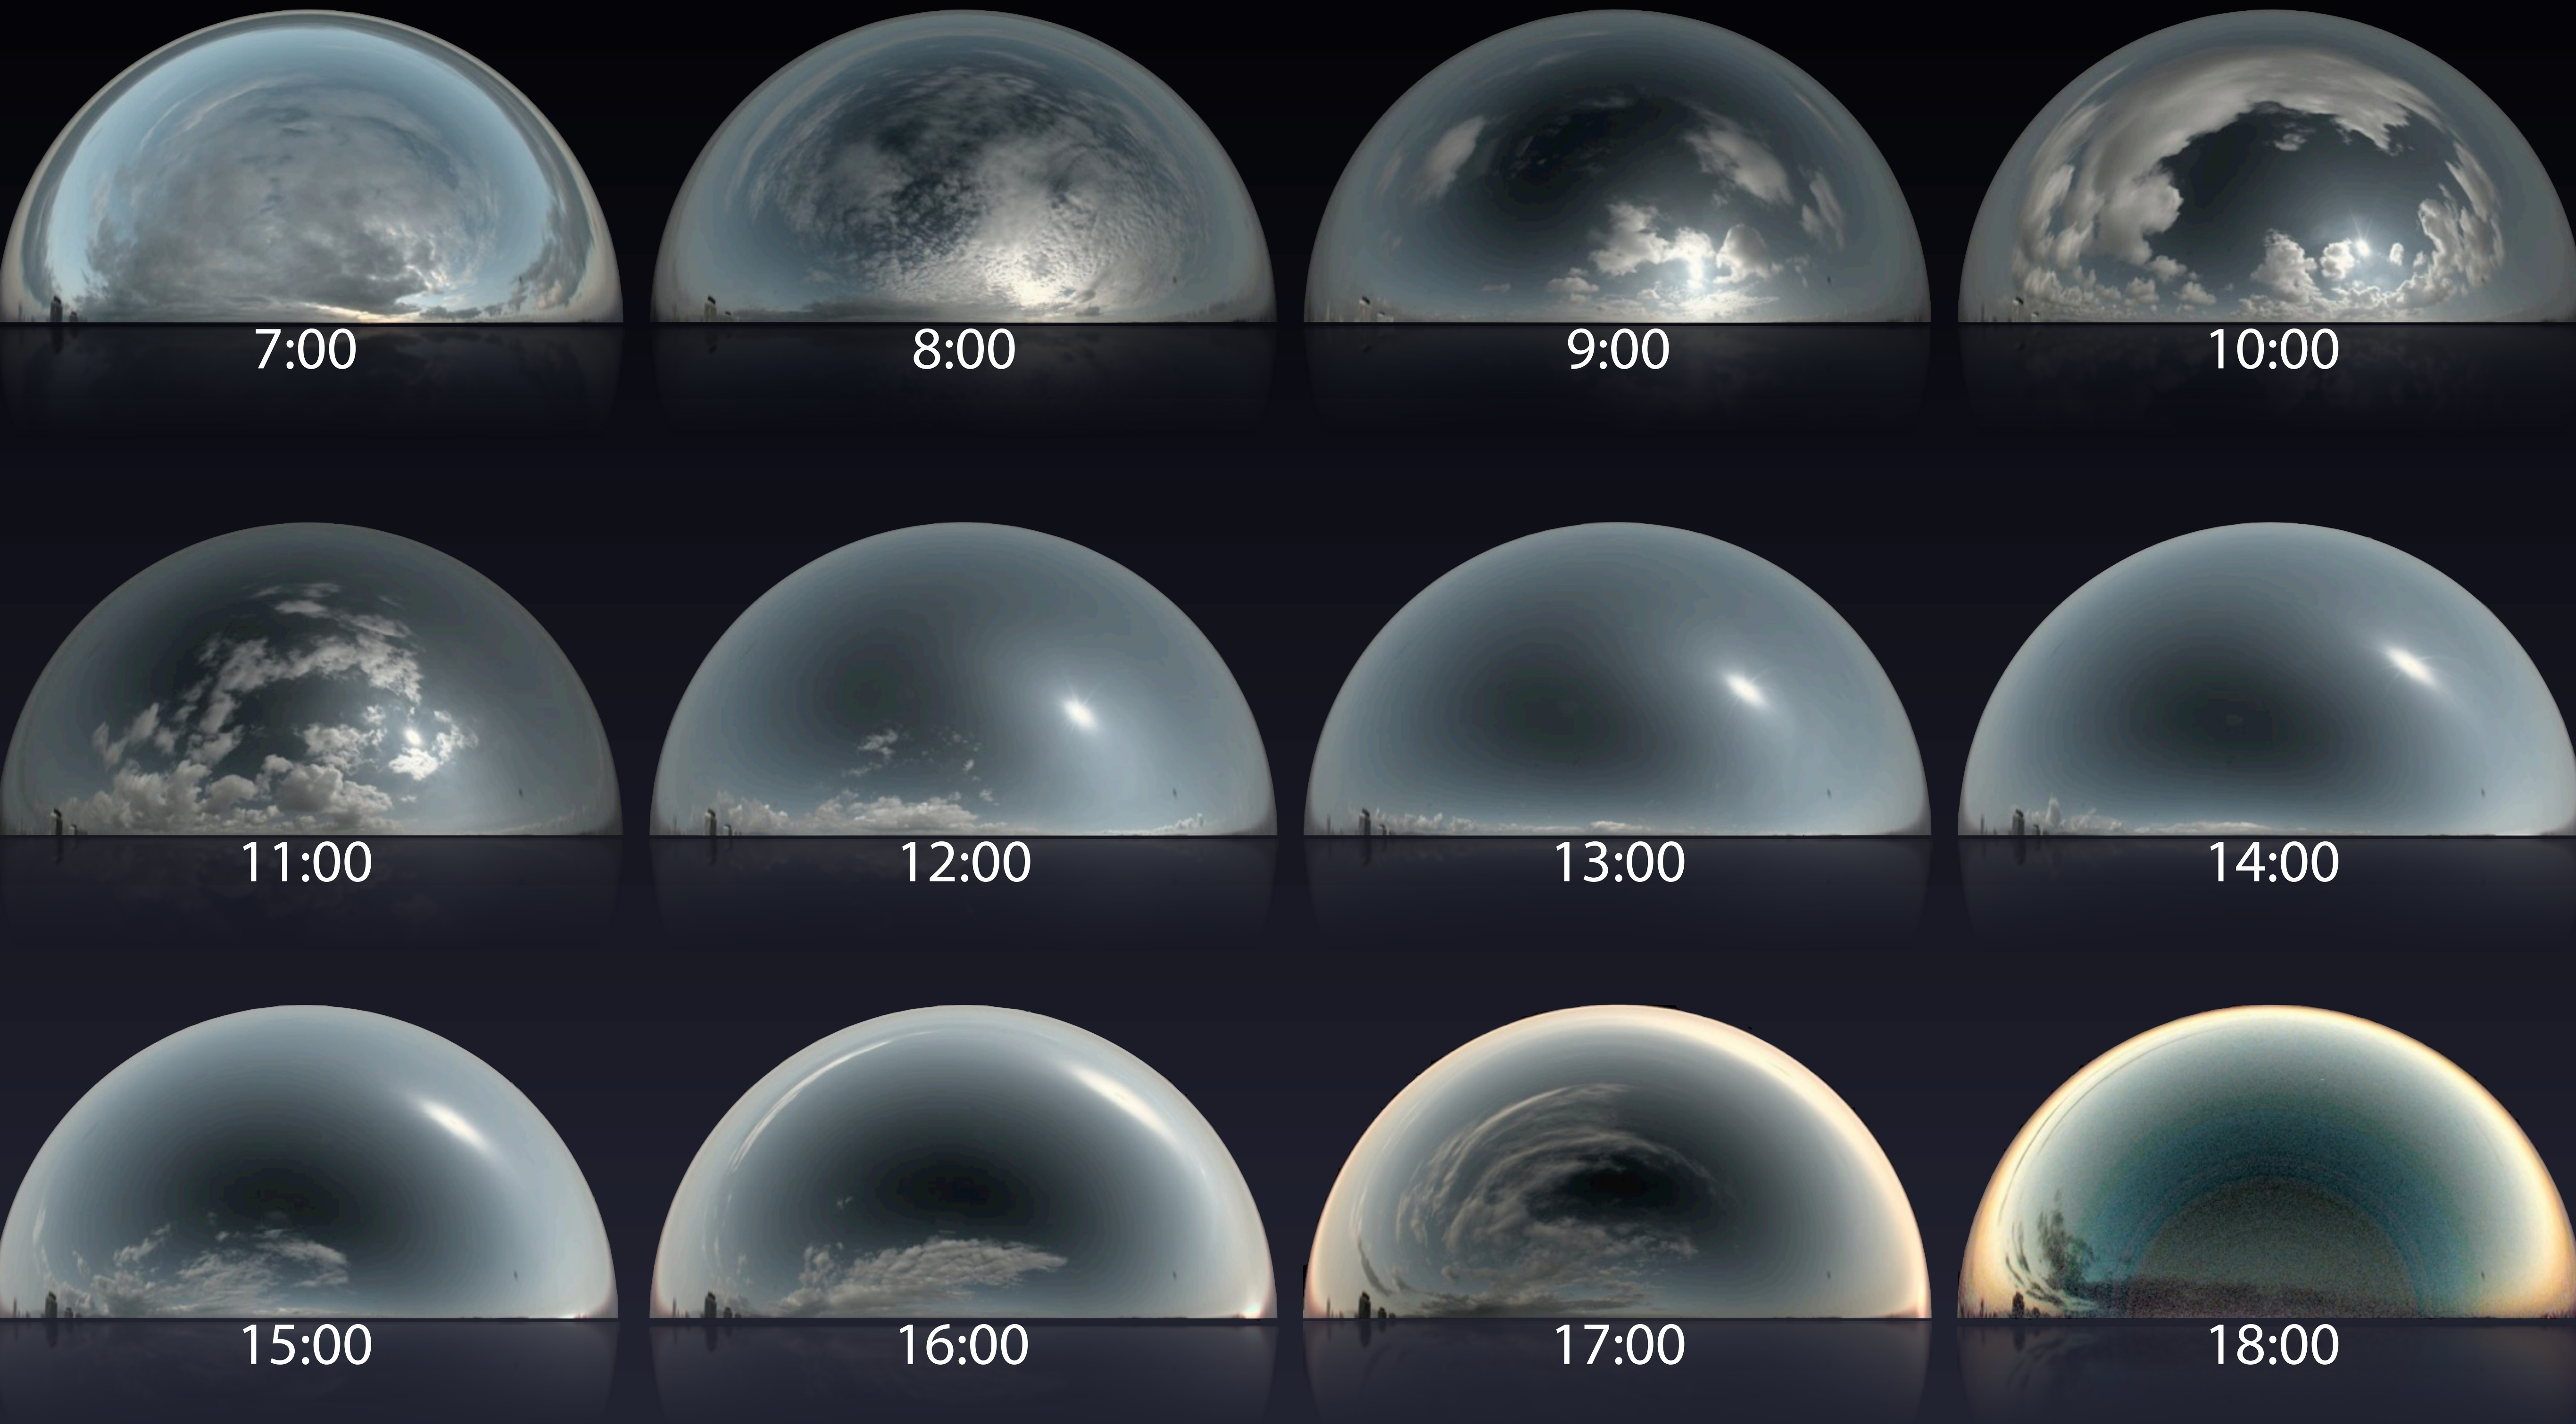

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Stumpfel et al., 2004]

So we adopt a data-driven approach, and use the sky database of Debevec et al, which contains HDR pictures of the sky hemisphere taken over an entire day. We detect the sun at each frame where it is visible, and extract a small window around it. Note that I'm showing the sun positions in image space for visualization, but in practice we perform the computations in angular space to avoid any distortions.

Environment map – sun

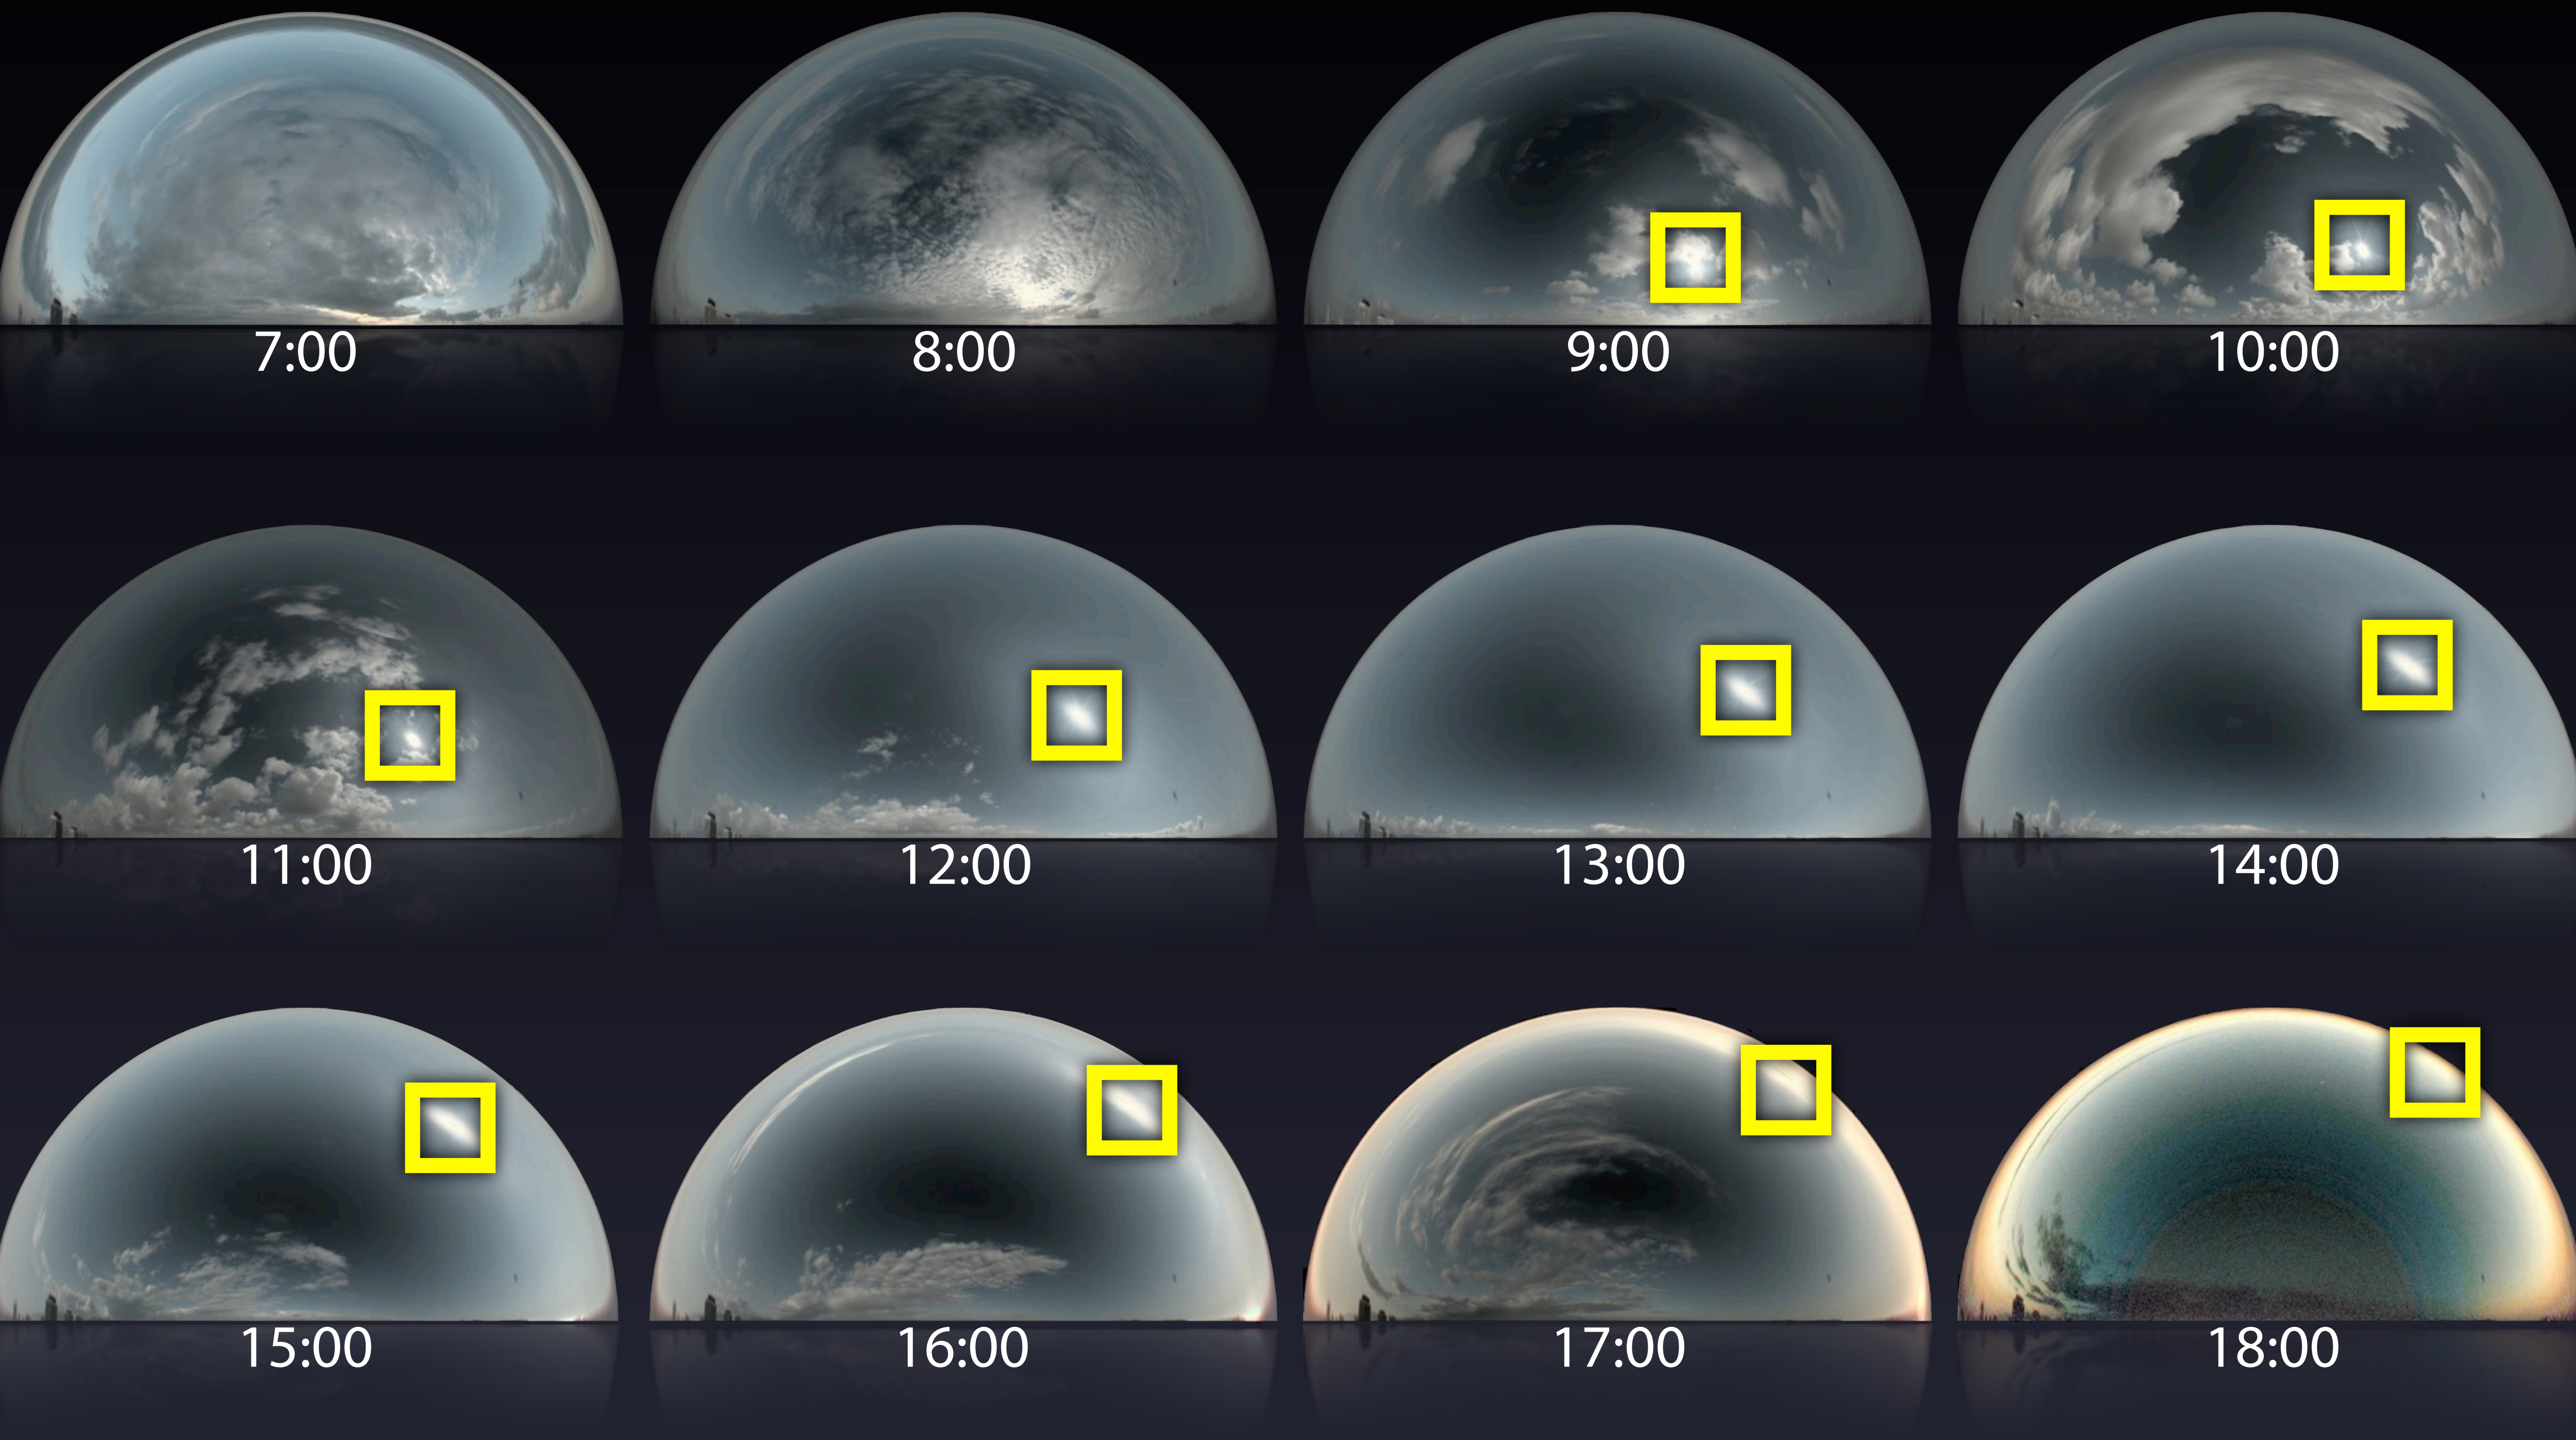

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Stumpfel et al., 2004]

So we adopt a data-driven approach, and use the sky database of Debevec et al, which contains HDR pictures of the sky hemisphere taken over an entire day. We detect the sun at each frame where it is visible, and extract a small window around it. Note that I'm showing the sun positions in image space for visualization, but in practice we perform the computations in angular space to avoid any distortions.

Environment map – sun

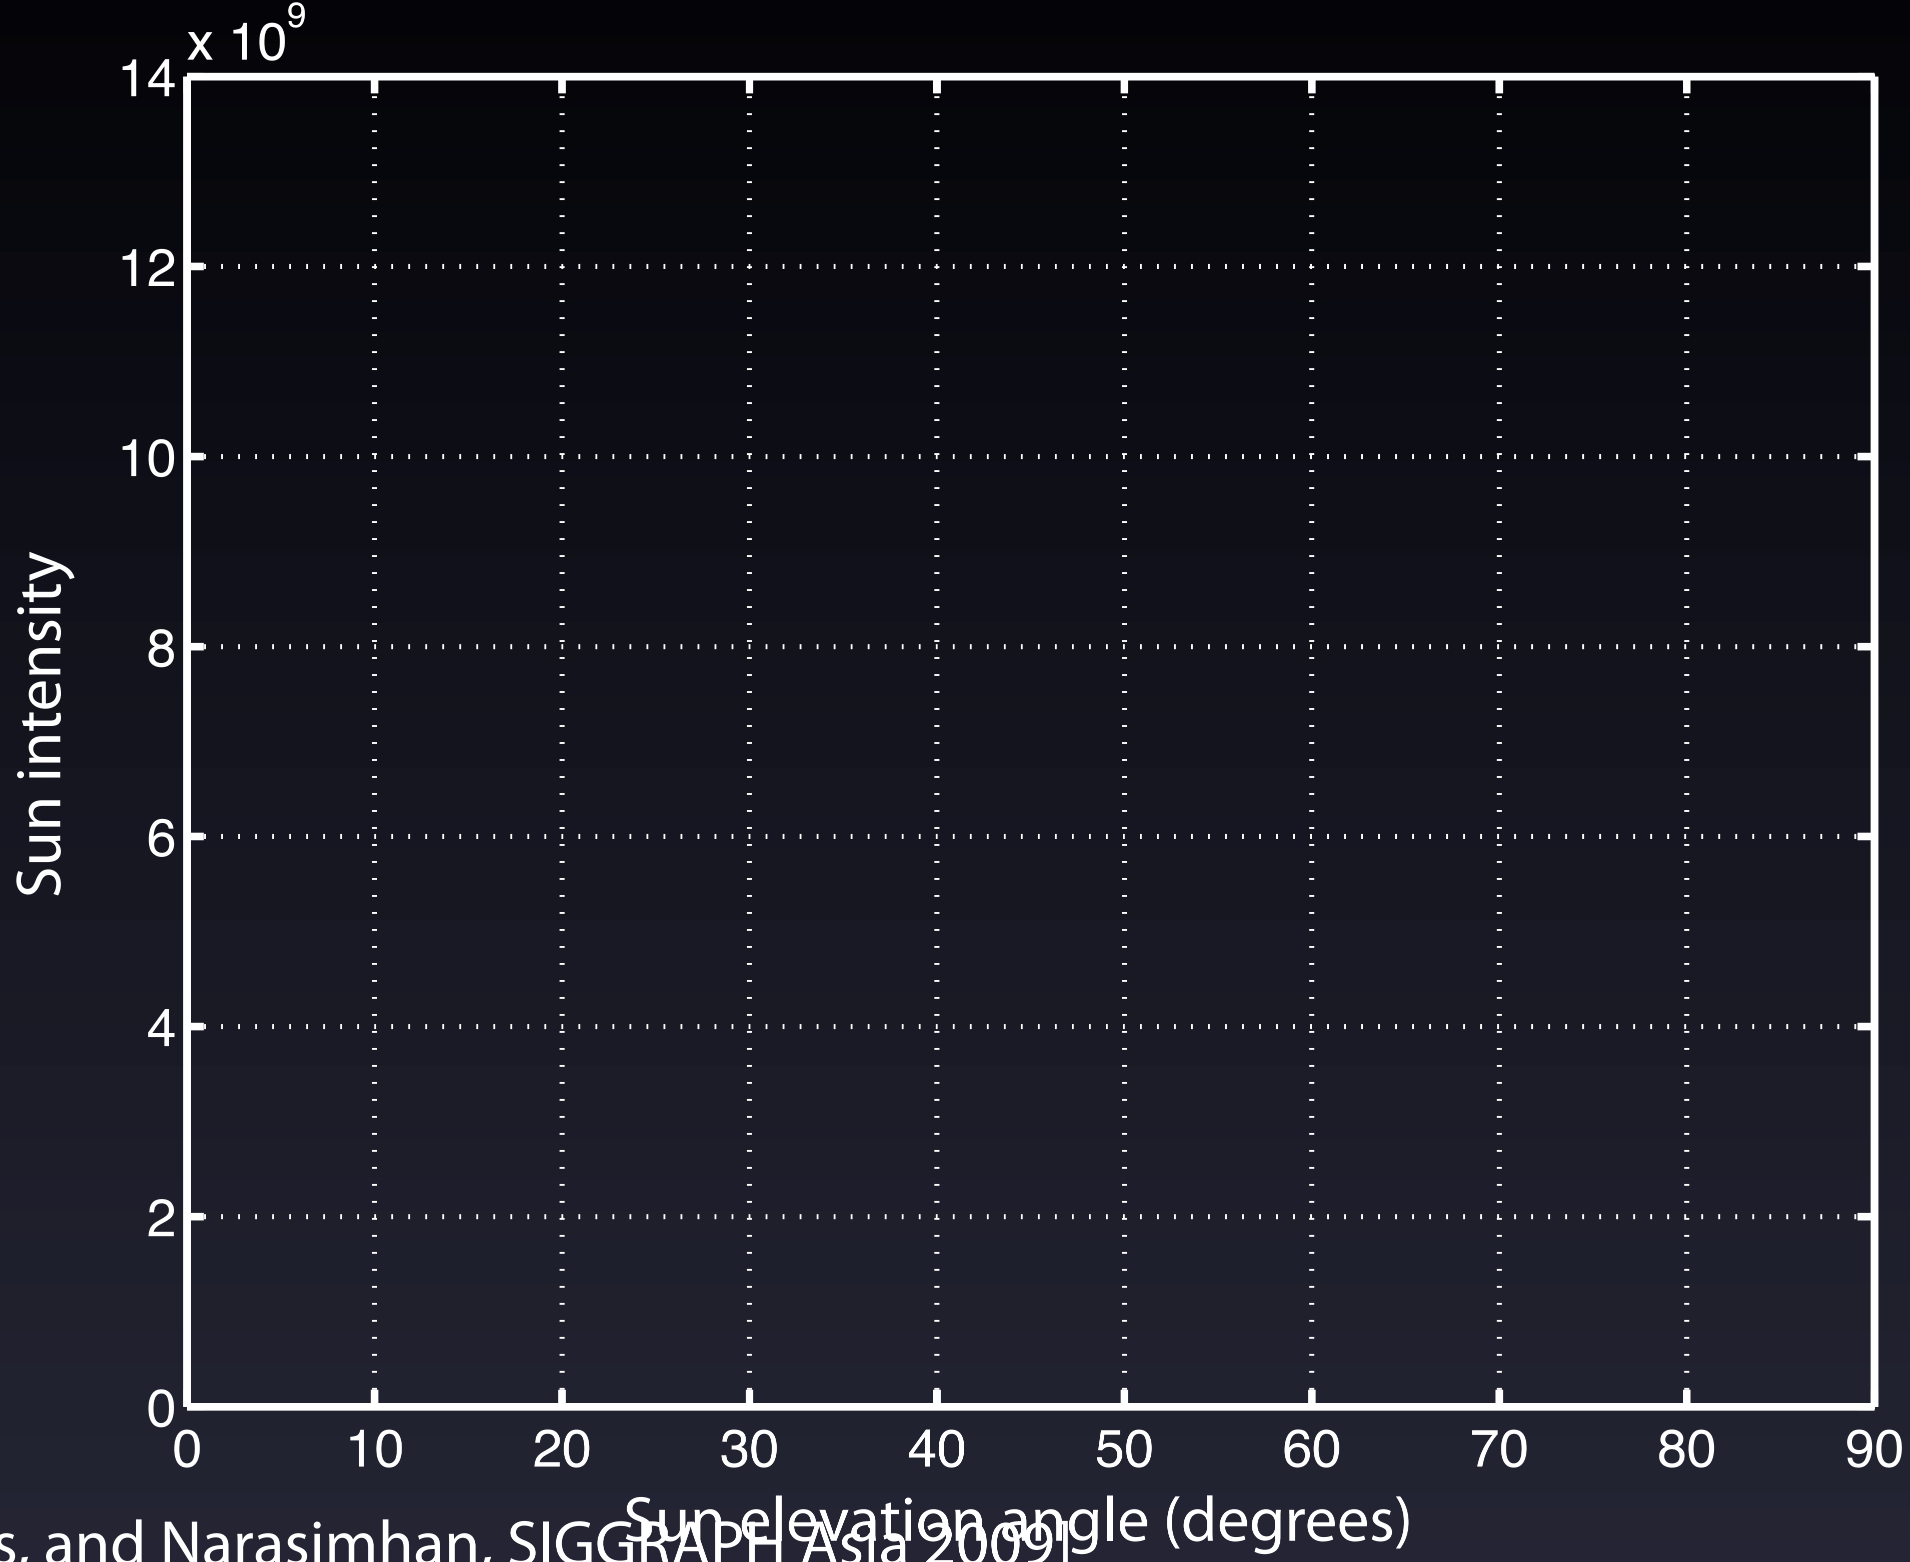

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We then plot the sun intensity at each image from the dataset as a function of elevation angle as extracted from the dataset, we can fit an atmospheric model to it and obtain a parametric, data-driven model of how the sun looks like for all elevation angles.
As you can see here, it models a red sun at low elevation angles, and becomes brighter and yellower when the sun is higher.
Note that our model appears saturated, but that's only because of the limitations of the 8-bits display, If we plot its appearance in log space, we see that we've captured the full dynamic range.

Environment map – sun

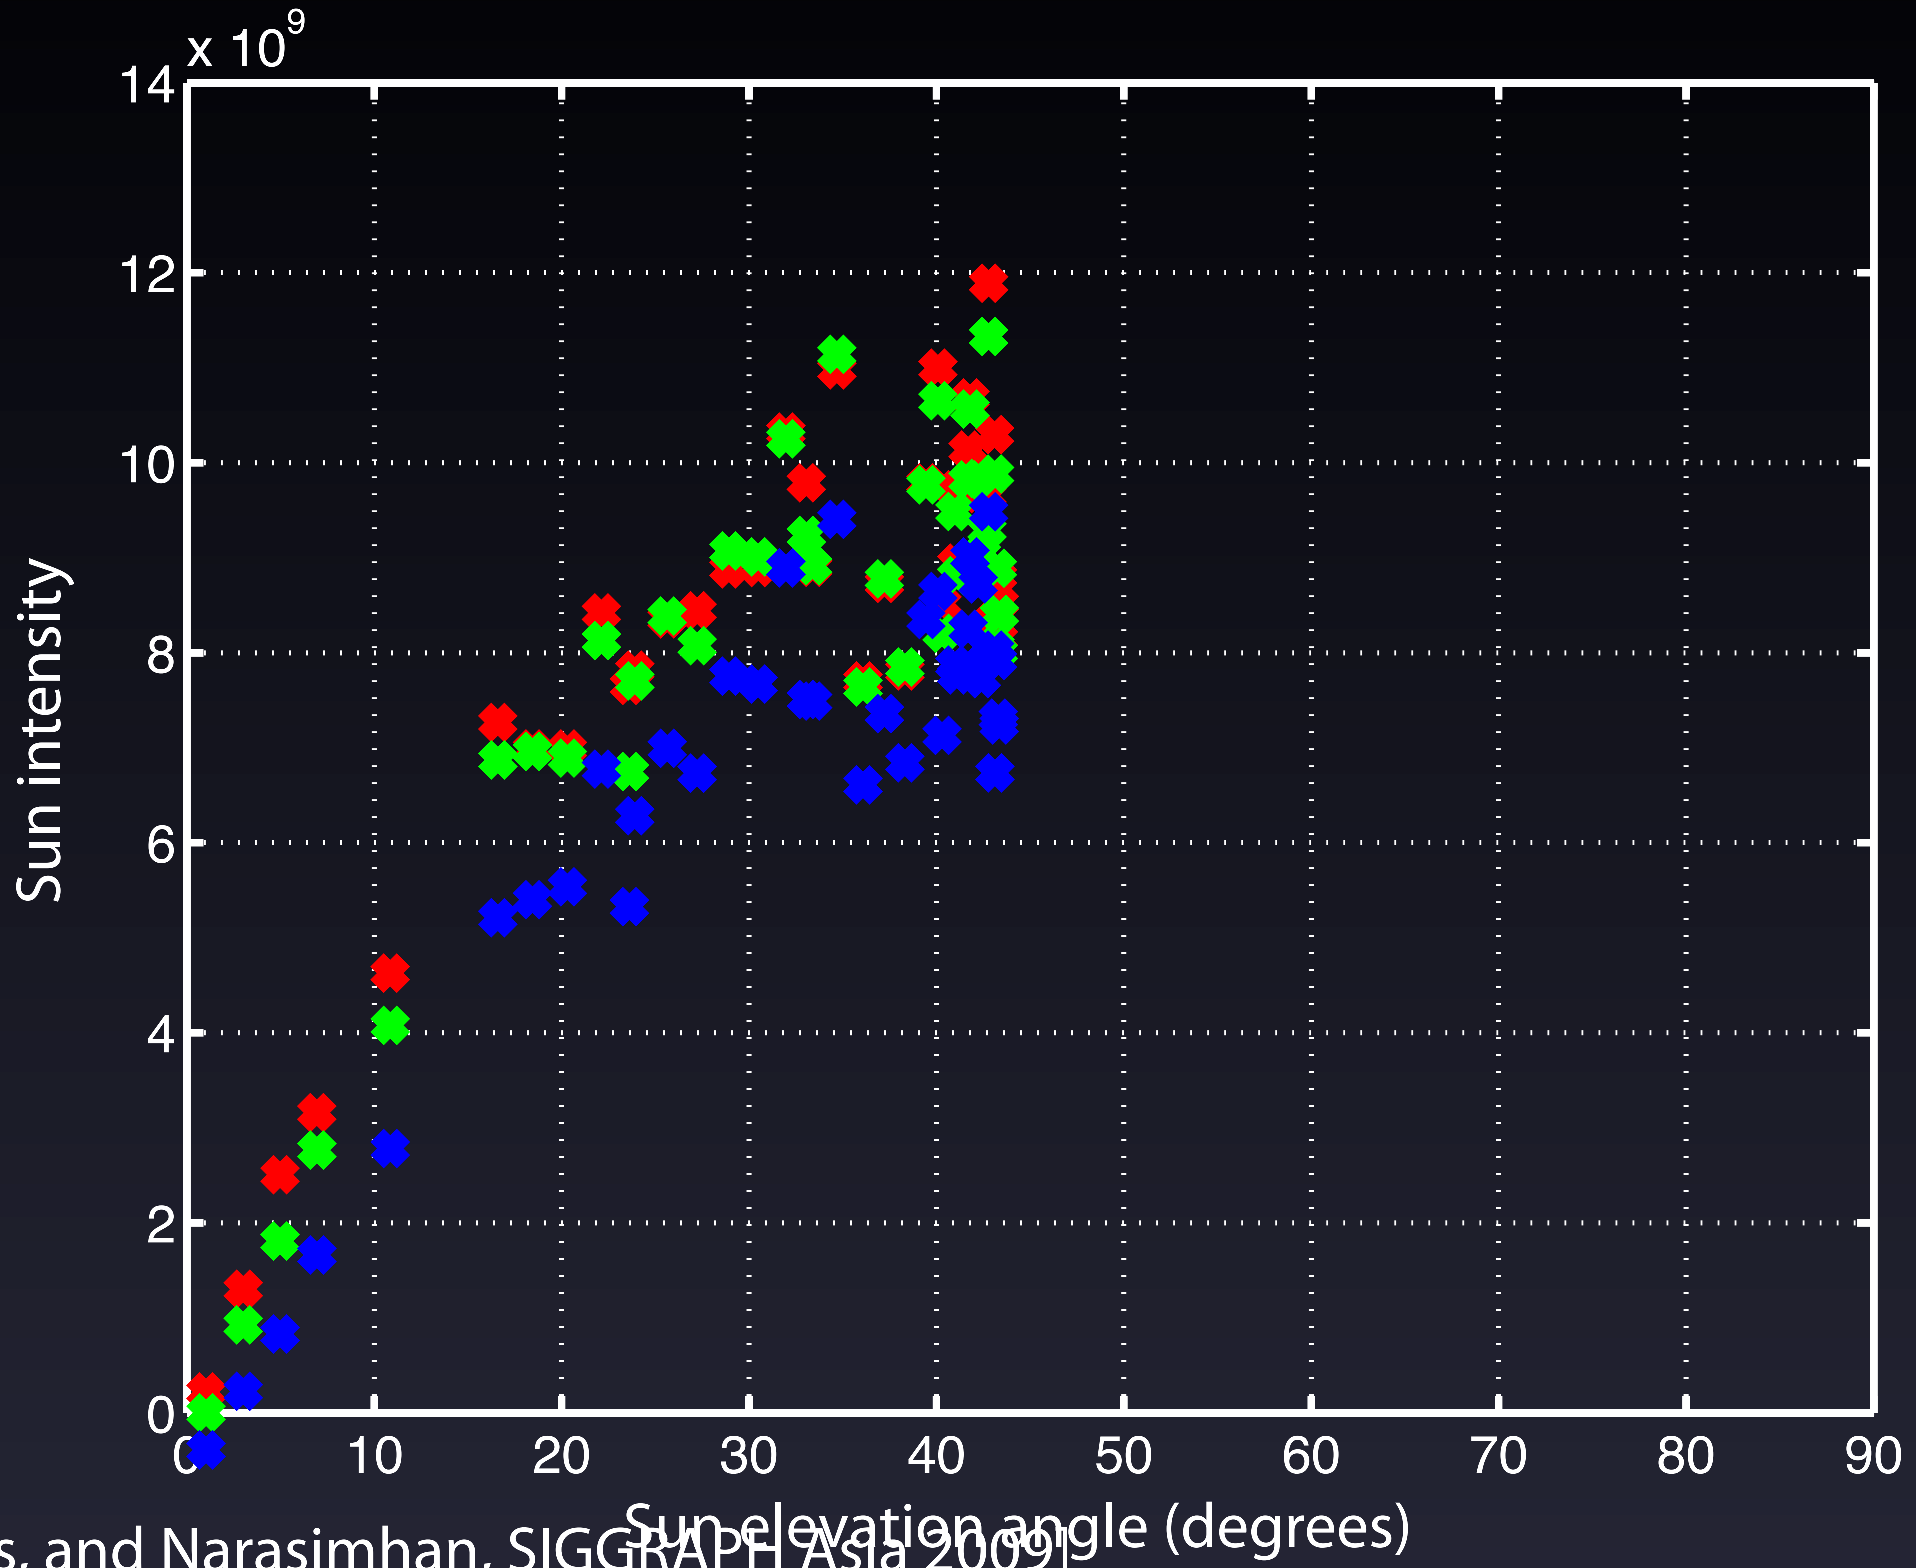

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We then plot the sun intensity at each image from the dataset as a function of elevation angle as extracted from the dataset, we can fit an atmospheric model to it and obtain a parametric, data-driven model of how the sun looks like for all elevation angles.
As you can see here, it models a red sun at low elevation angles, and becomes brighter and yellower when the sun is higher.
Note that our model appears saturated, but that's only because of the limitations of the 8-bits display, If we plot its appearance in log space, we see that we've captured the full dynamic range.

Environment map – sun

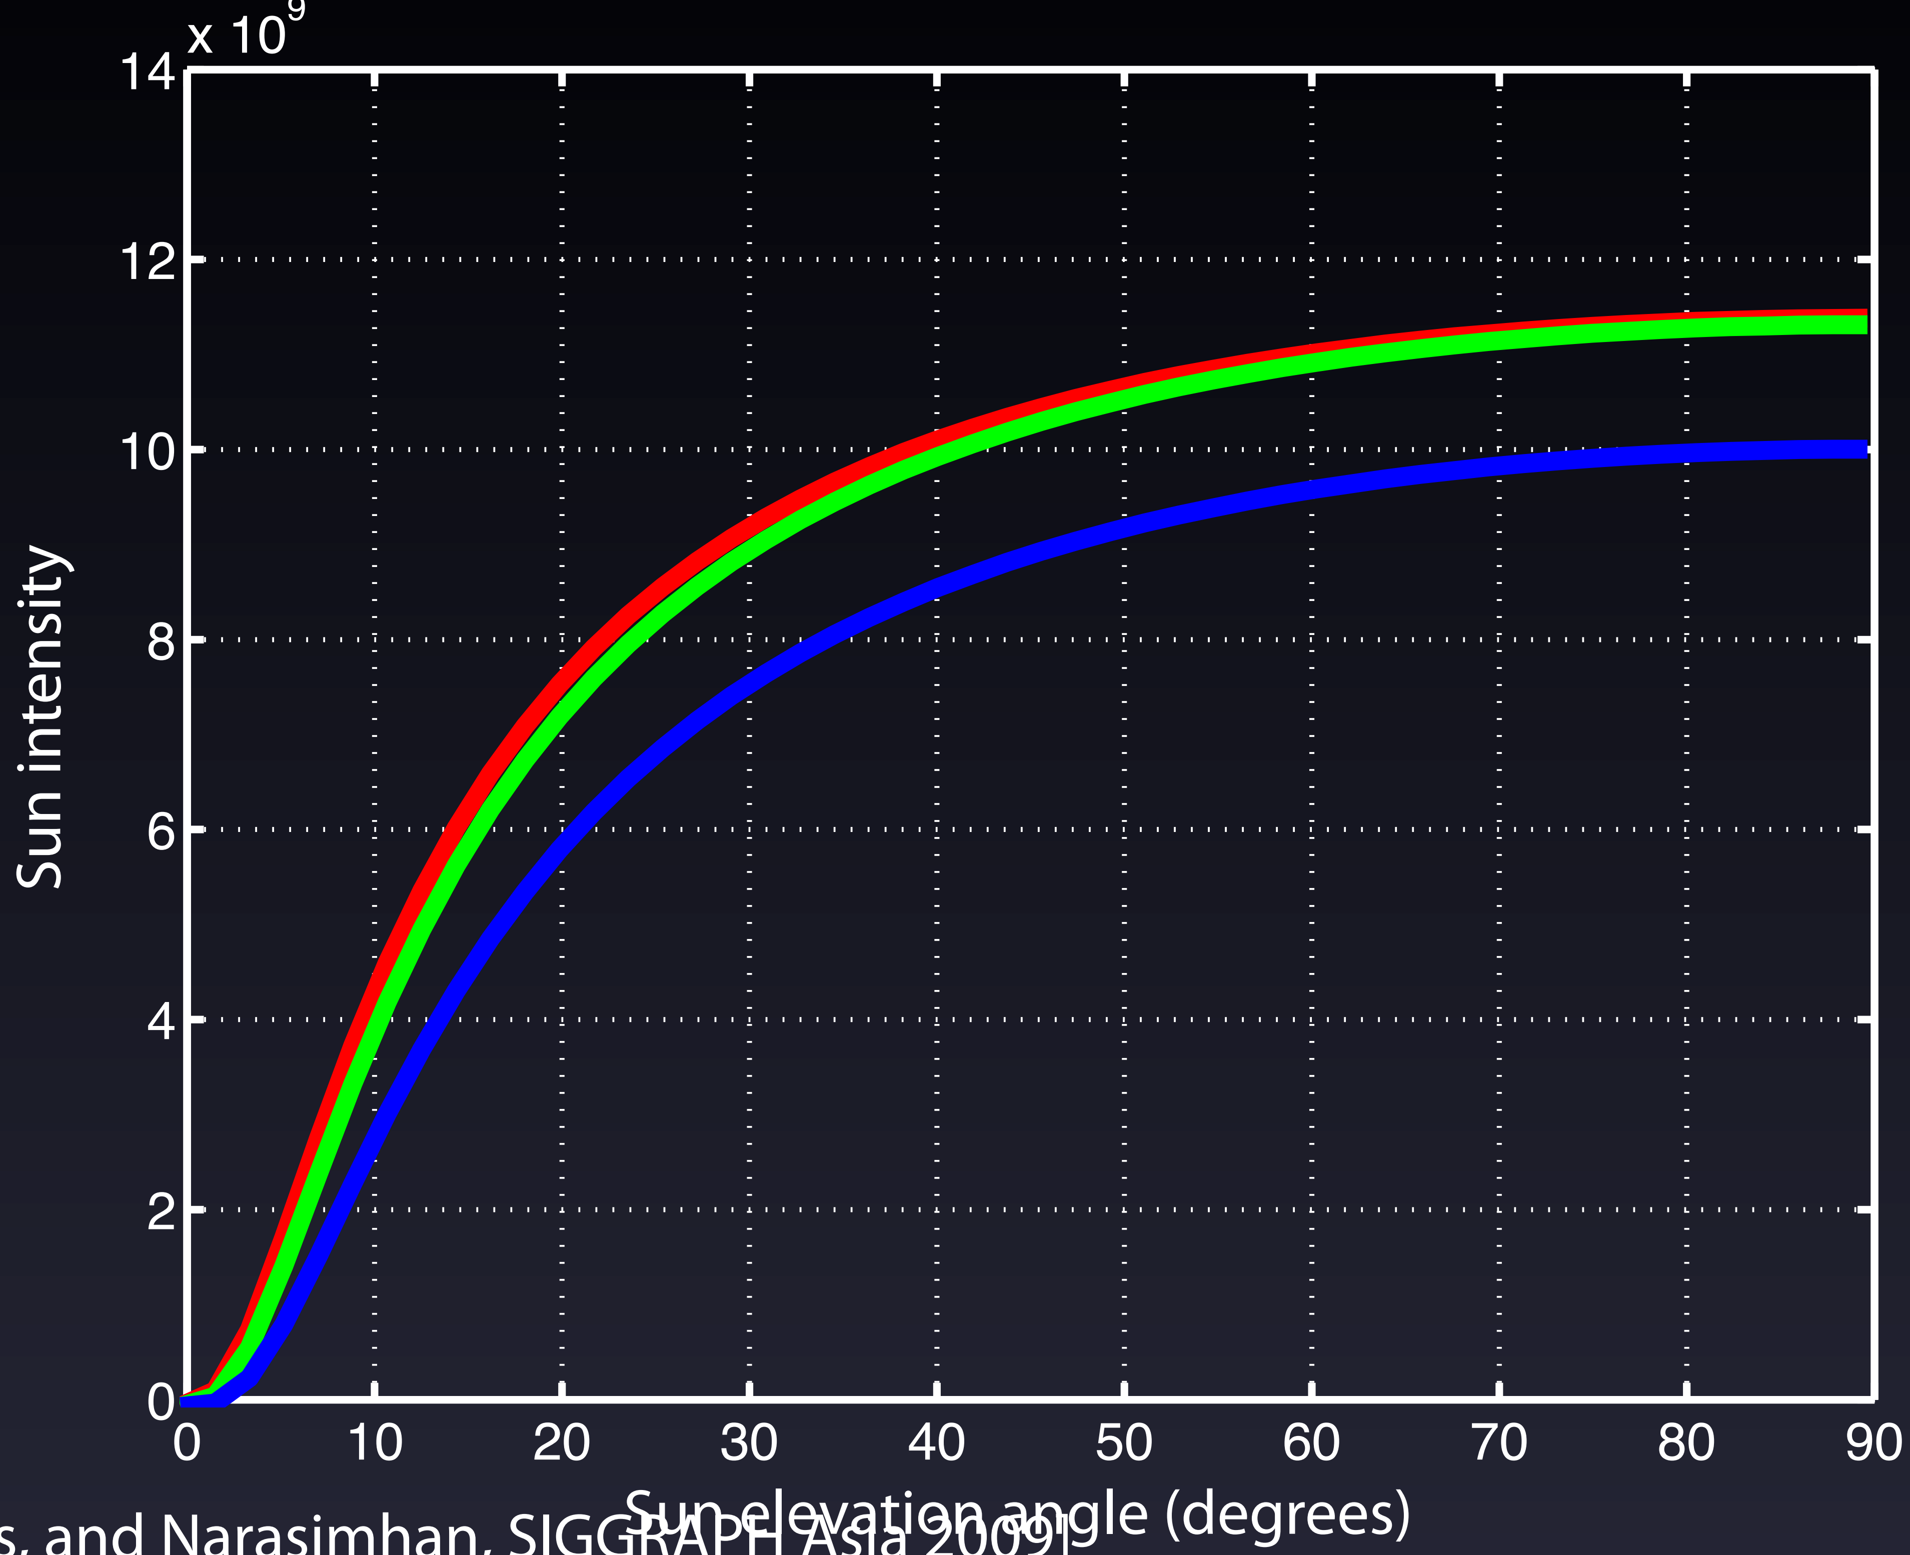

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We then plot the sun intensity at each image from the dataset as a function of elevation angle as extracted from the dataset, we can fit an atmospheric model to it and obtain a parametric, data-driven model of how the sun looks like for all elevation angles.
As you can see here, it models a red sun at low elevation angles, and becomes brighter and yellower when the sun is higher.
Note that our model appears saturated, but that’s only because of the limitations of the 8-bits display, If we plot its appearance in log space, we see that we’ve captured the full dynamic range.

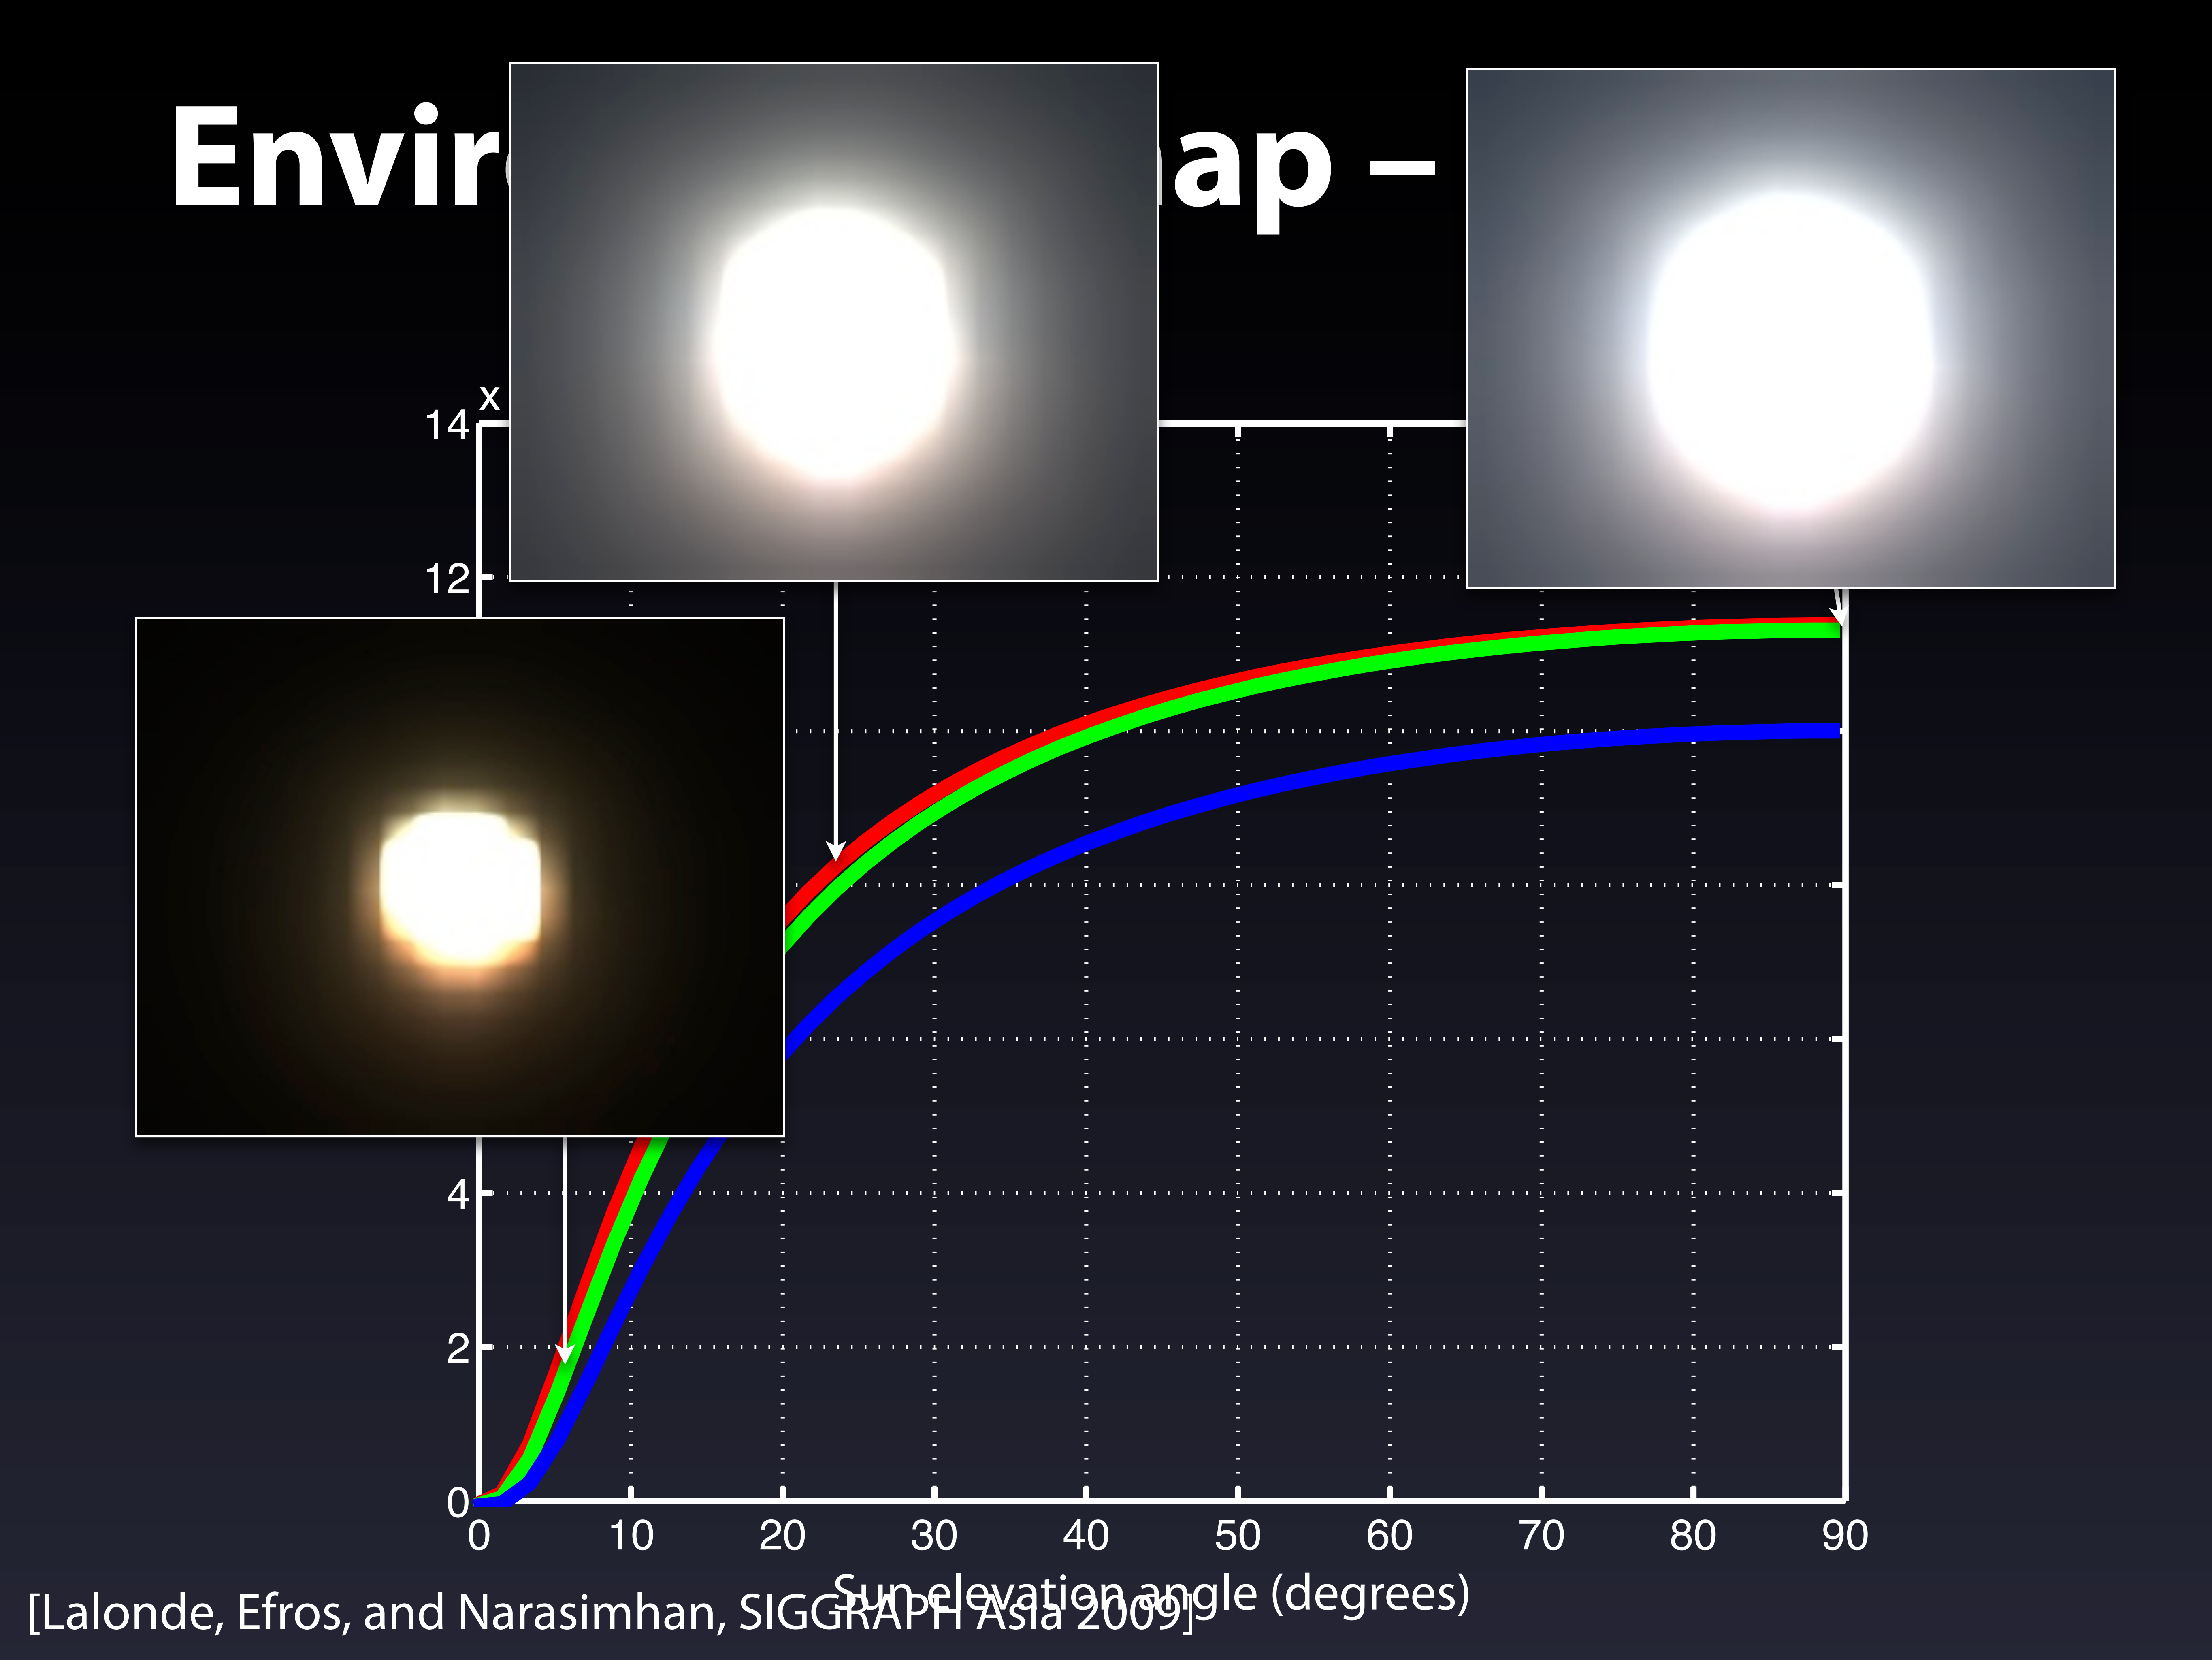

We then plot the sun intensity at each image from the dataset as a function of elevation angle as extracted from the dataset, we can fit an atmospheric model to it and obtain a parametric, data-driven model of how the sun looks like for all elevation angles.
As you can see here, it models a red sun at low elevation angles, and becomes brighter and yellower when the sun is higher.
Note that our model appears saturated, but that's only because of the limitations of the 8-bits display, If we plot its appearance in log space, we see that we've captured the full dynamic range.

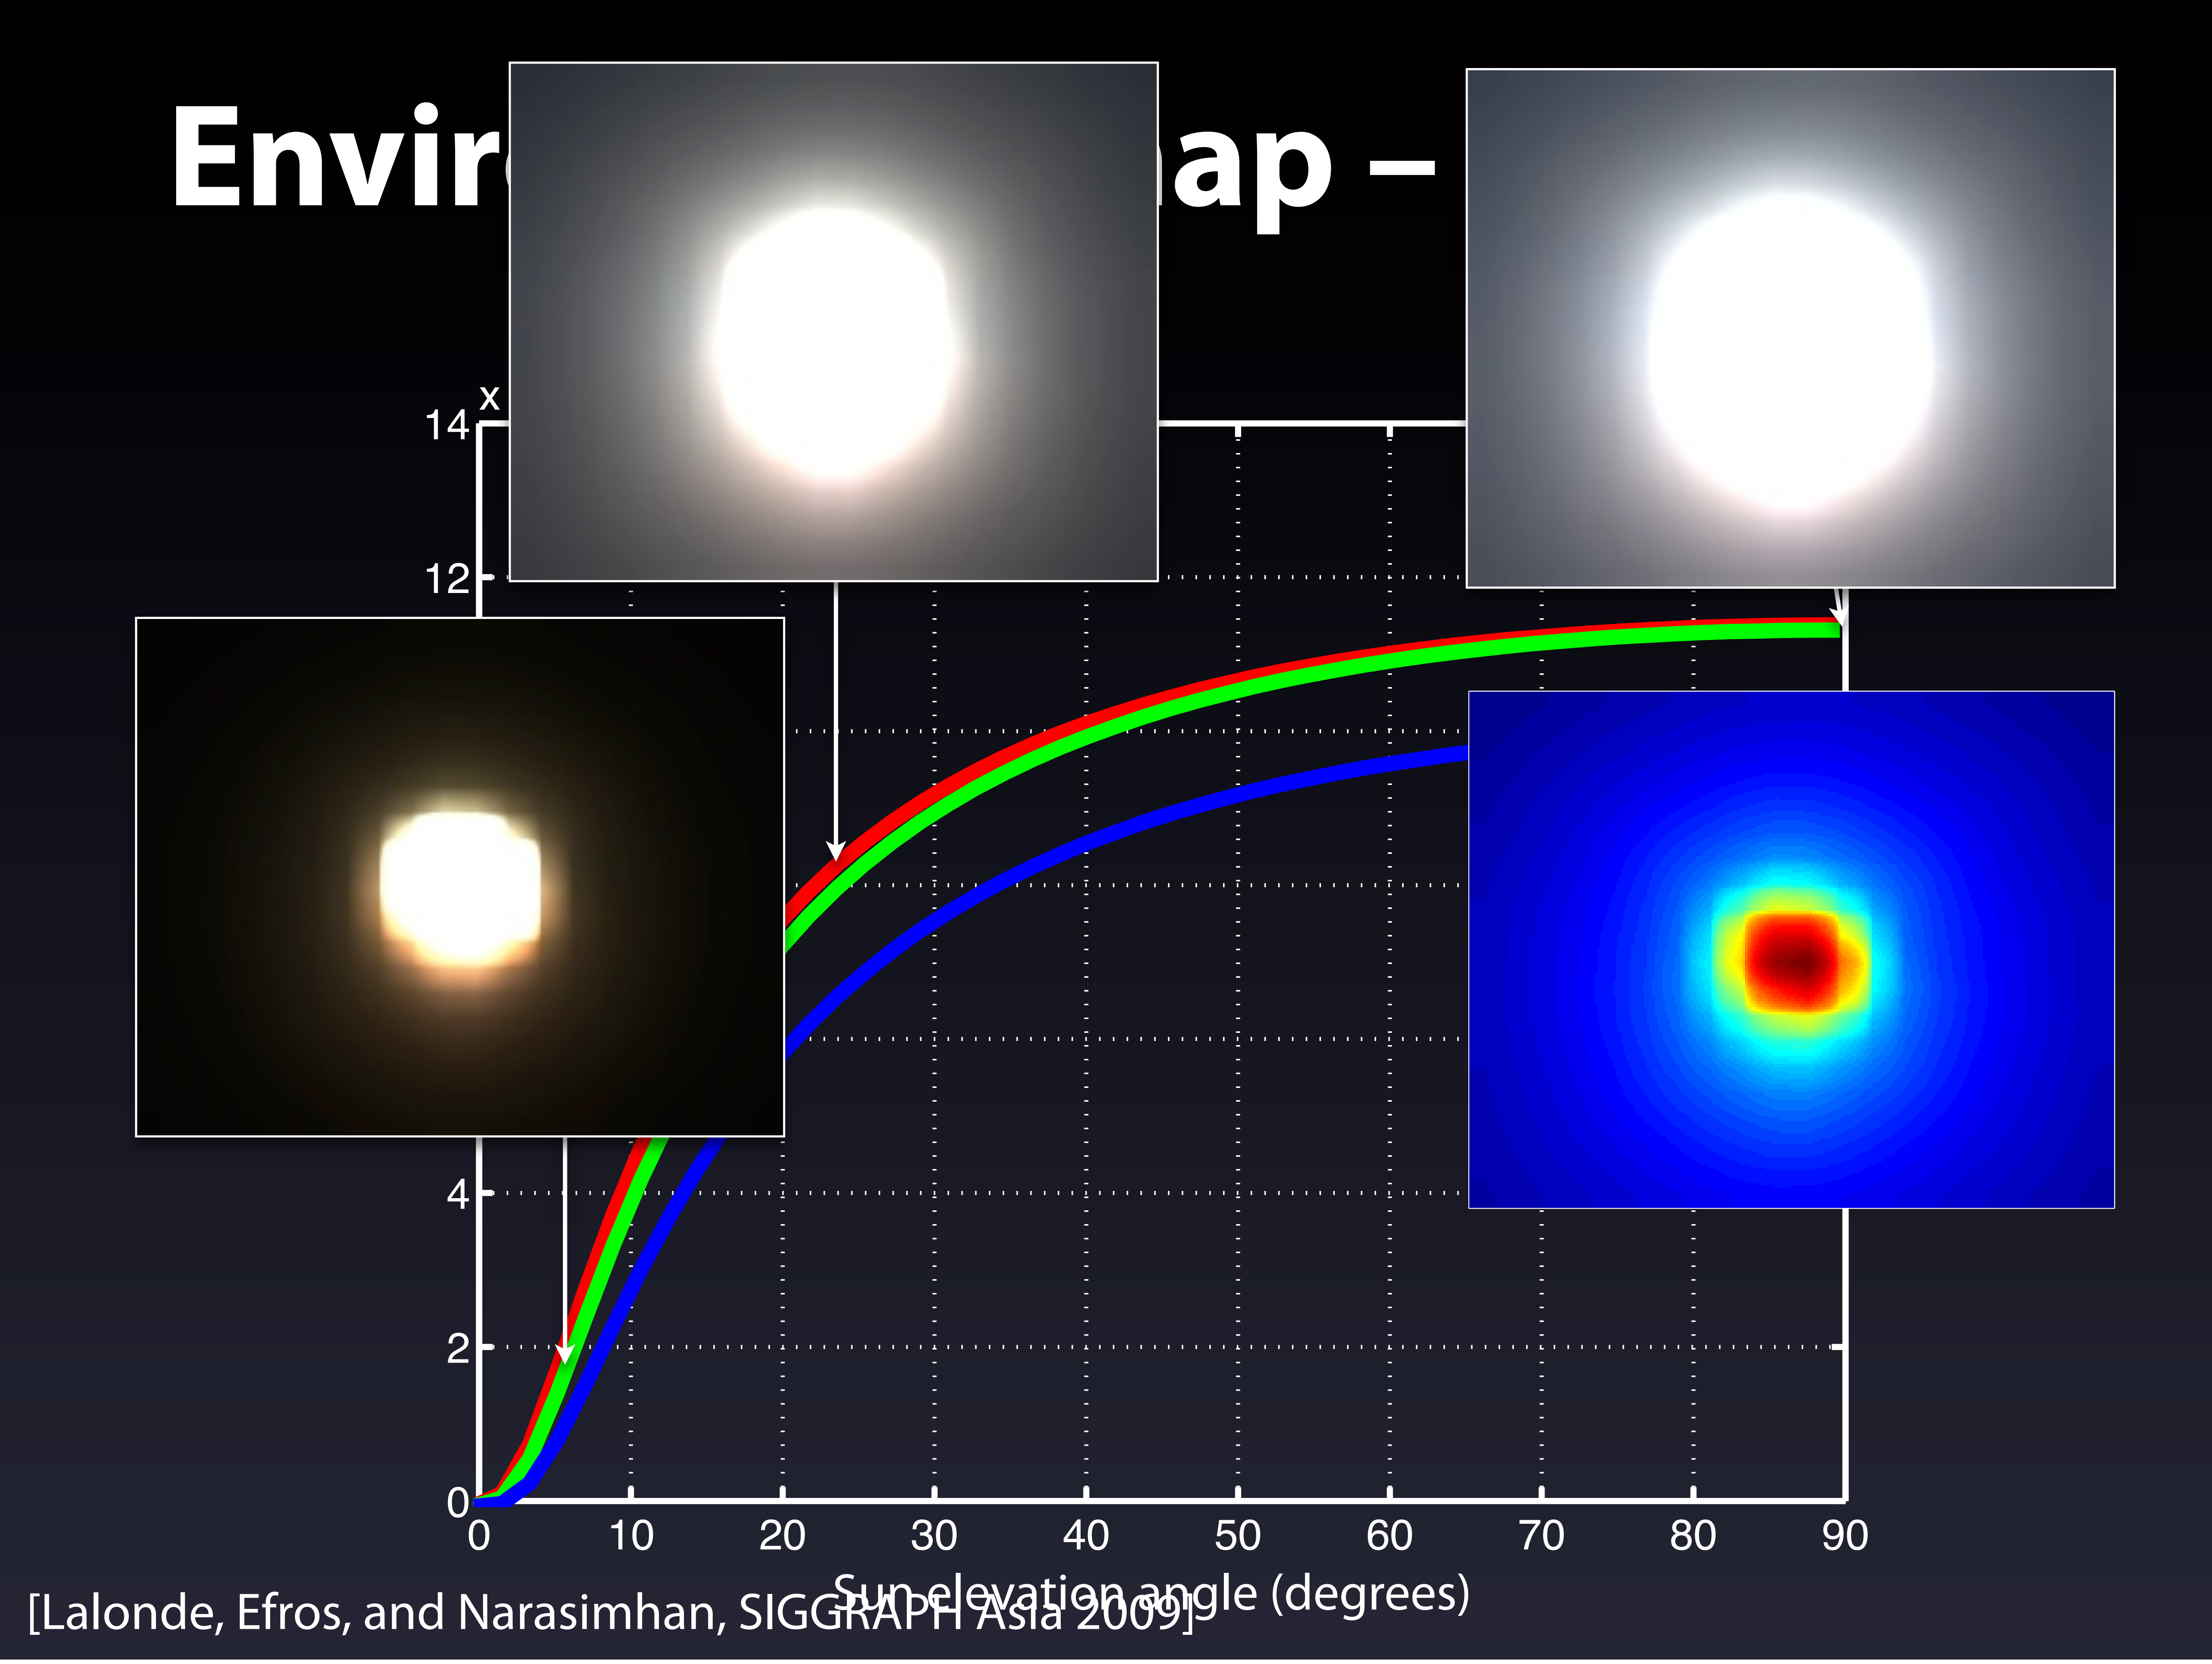

We then plot the sun intensity at each image from the dataset as a function of elevation angle as extracted from the dataset, we can fit an atmospheric model to it and obtain a parametric, data-driven model of how the sun looks like for all elevation angles.
As you can see here, it models a red sun at low elevation angles, and becomes brighter and yellower when the sun is higher.
Note that our model appears saturated, but that’s only because of the limitations of the 8-bits display, If we plot its appearance in log space, we see that we’ve captured the full dynamic range.

Environment map – scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can’t see what’s behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we’ve already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – scene

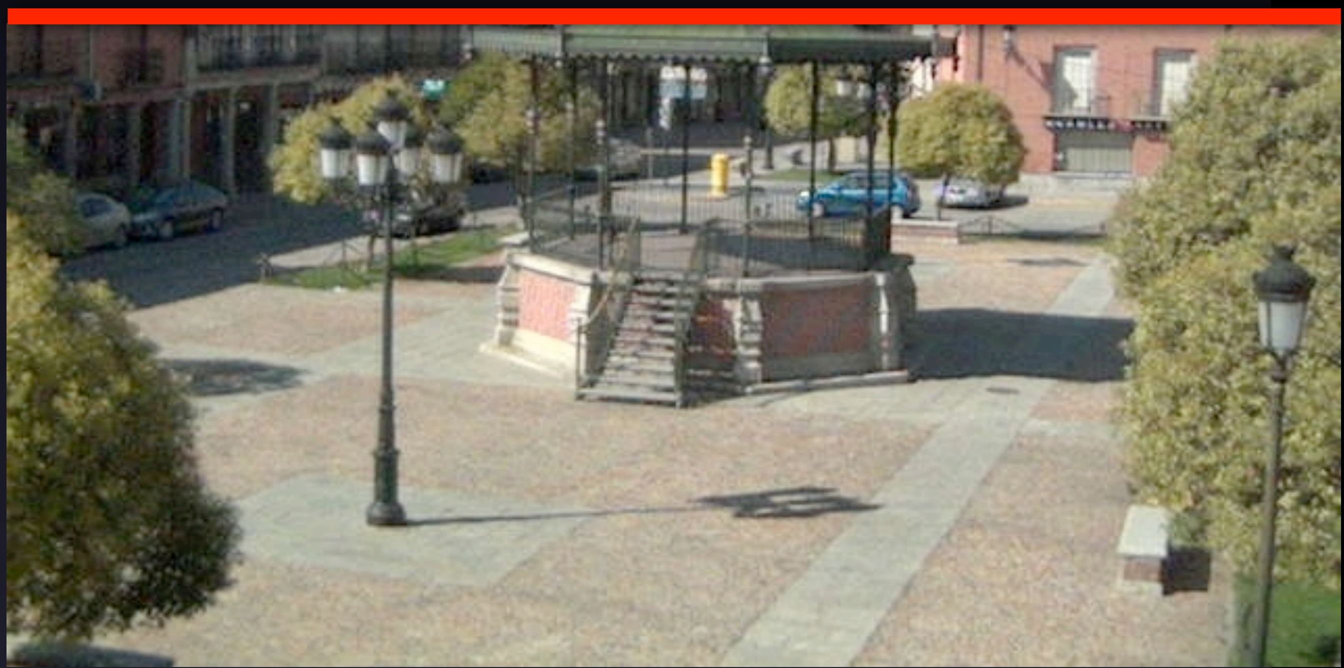

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – scene

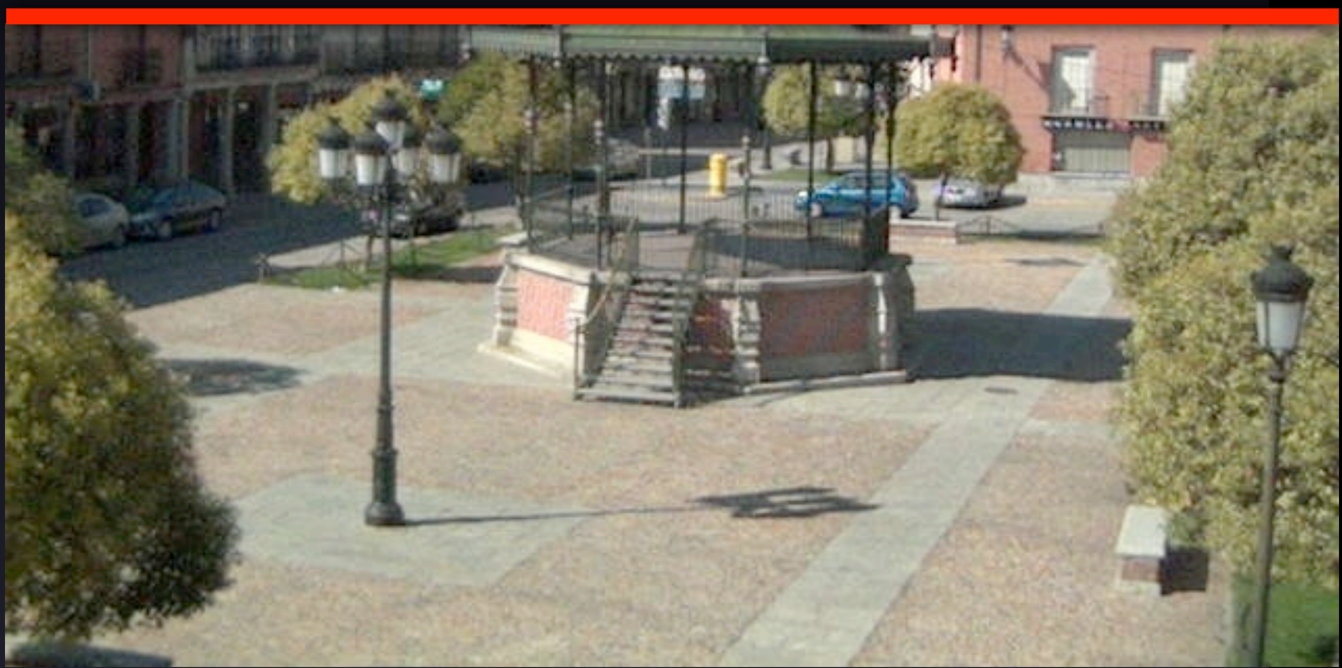

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Khan *et al.*, 2006]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

[Khan *et al.*, 2006]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.

Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – clouds

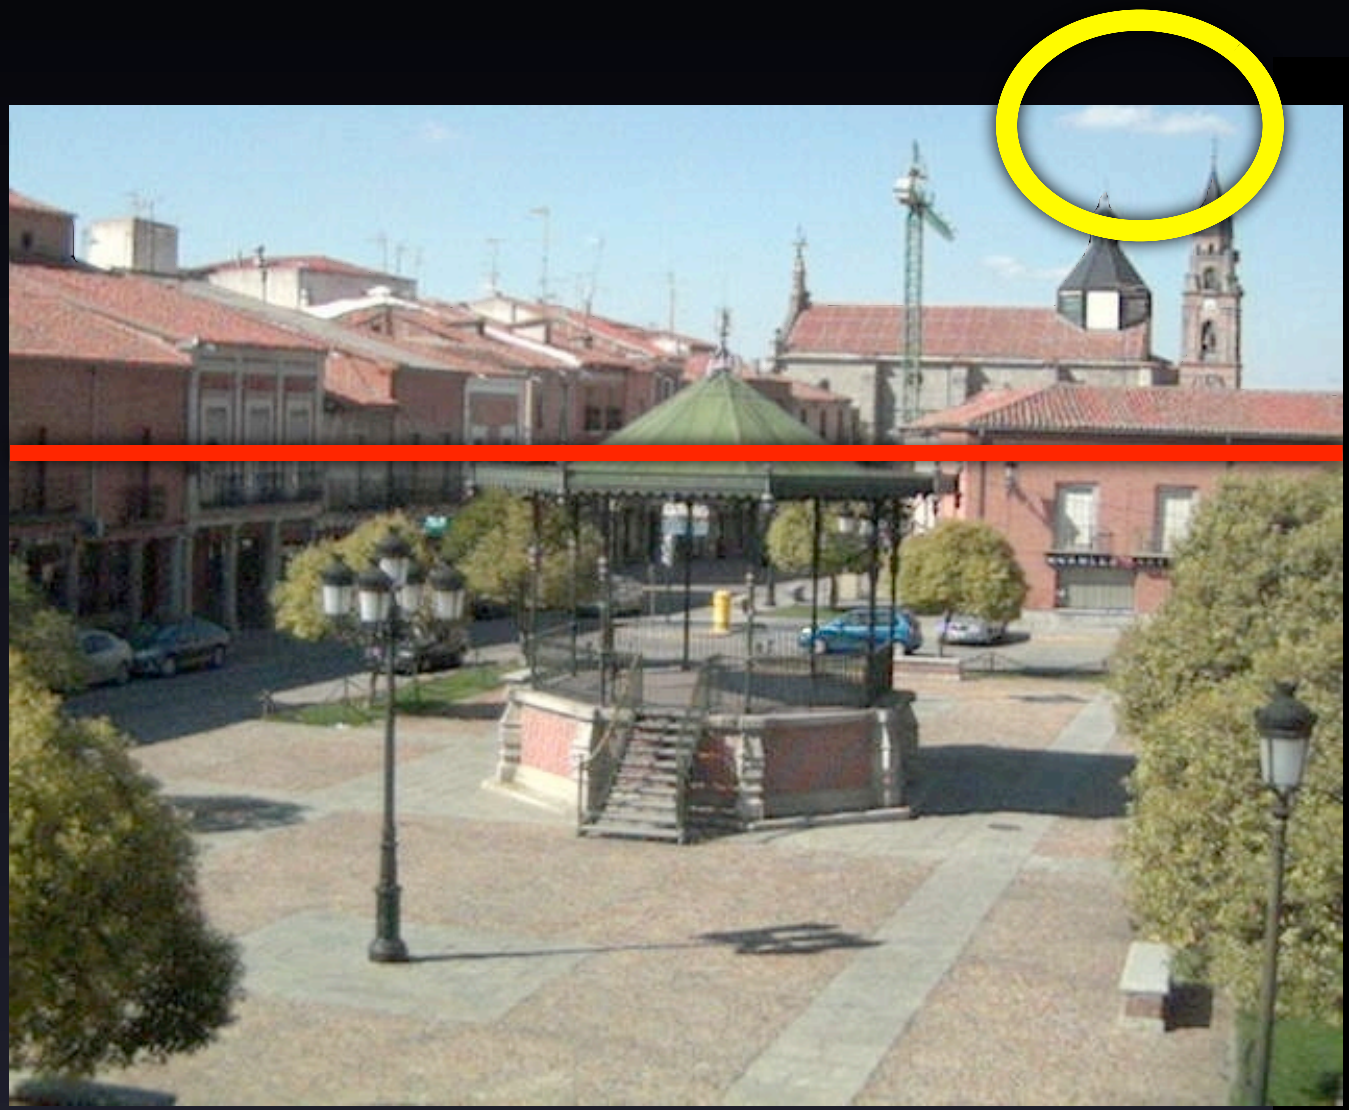

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

Environment map – clouds

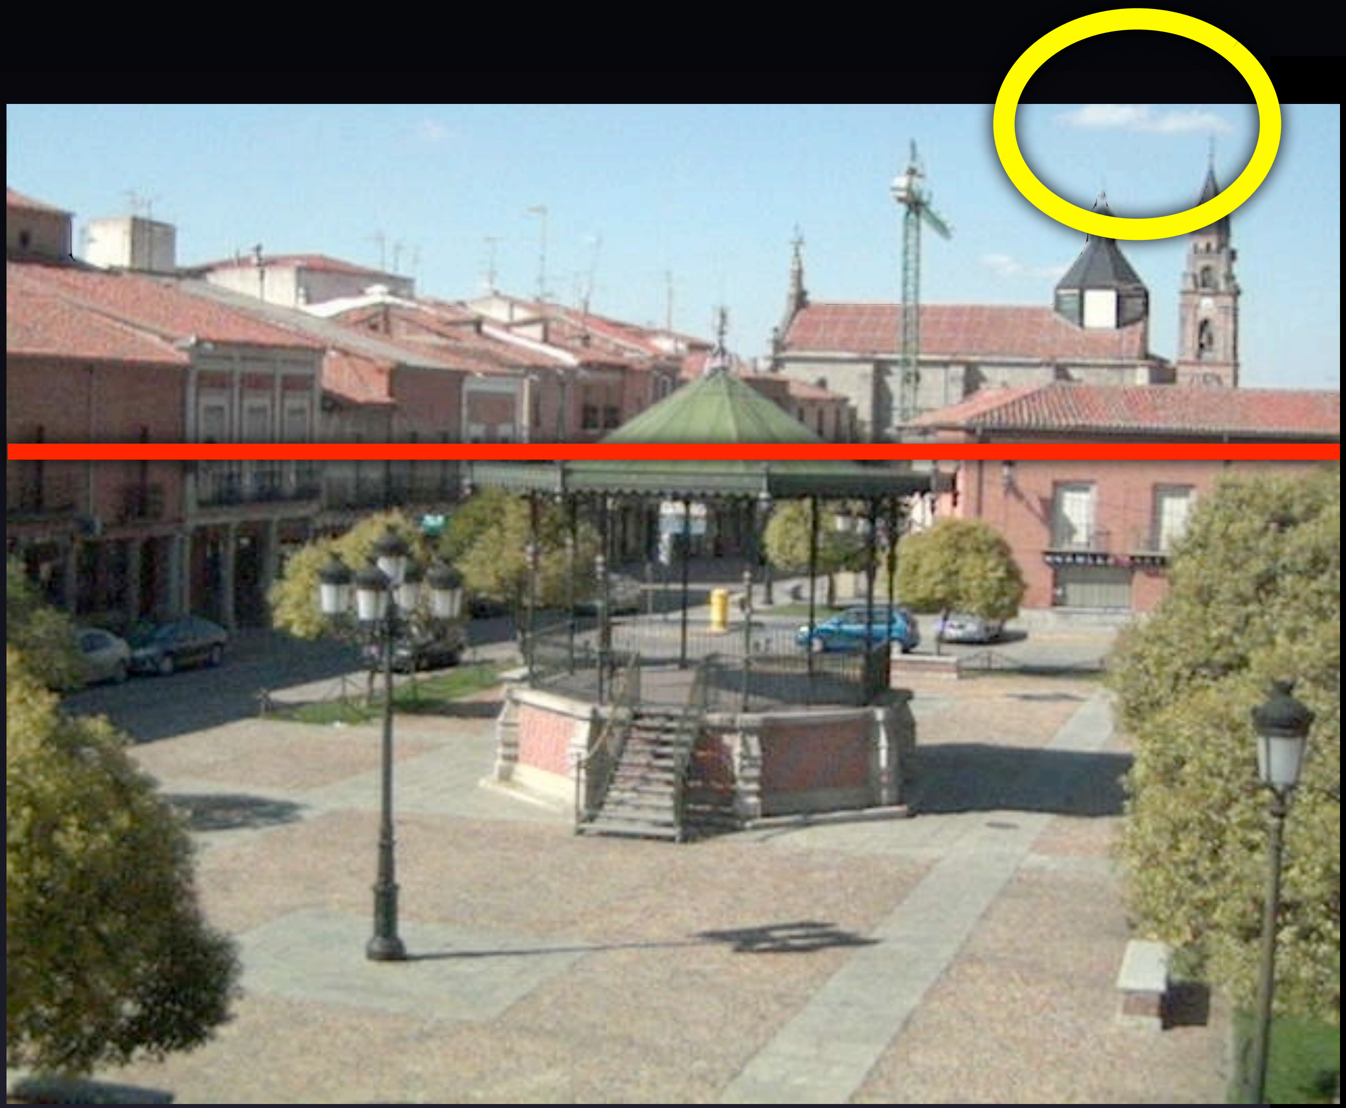

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This gives us a nice physics-based, HDR sky hemisphere. For the scene, we can't see what's behind us, so we use the approximation that was proposed by Khan et al. in 2006, and project what lies below the horizon onto the bottom hemisphere of the environment map. For the parts of the scene which protrude above the horizon, and occlude the sky, we use the same trick, except that we make sure that the highest point in the scene gets projected to the right elevation angle.
Finally, the last missing piece is the clouds. For this we use texture synthesis and replicate the clouds that we've already segmented. This does not yield physically-correct results, but since we already know whether the sun is visible or not, we need clouds only to generate believable reflections off of shiny objects.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We can then use our environment maps to relight our 3-D model of the statue, and realistically insert it into the scene. Notice how the shadows are consistent with the image, and how they also disappear when the sun does not shine on the scene.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We can then use our environment maps to relight our 3-D model of the statue, and realistically insert it into the scene. Notice how the shadows are consistent with the image, and how they also disappear when the sun does not shine on the scene.

[Khan et al., 2006]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We compare our method with the approach of Khan et al, which wraps the entire image around on a sphere. So here are 2 cars added to 2 frames from a webcam sequence. As you can see, the results don’t look too convincing. If we relight the same cars with our environment maps, it shows the importance of capturing the high dynamic range of outdoor lighting and the accurate direction of sun if we want to make the inserted objects realistic.

We compare our method with the approach of Khan et al, which wraps the entire image around on a sphere. So here are 2 cars added to 2 frames from a webcam sequence. As you can see, the results don’t look too convincing. If we relight the same cars with our environment maps, it shows the importance of capturing the high dynamic range of outdoor lighting and the accurate direction of sun if we want to make the inserted objects realistic.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

21

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there.
And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

This technique could also be useful to an architect who wishes to visualize how the house he just designed would look in natural settings.
For instance, all we need is a webcam of a beach, we can see how it would look as a beachfront property. Or maybe we also have a webcam of the countryside, we can see how it would look there. And how about in the Swiss alps? But the cool thing is: not only can we see it at different locations, but also at different times of day!
How does the beach look in broad sunlight? How does the countryside look at dusk? Where is the sun at 4pm in the mountains? These are just a few examples of the rich illumination effects that we're able to capture with our technique.

Capturing the illumination of the world

So far, we've been analyzing the illumination conditions of a few scenes only. But what we really want is to capture the illumination all around the world, and this is not just important for 3-d object insertion, but also other research areas in graphics, computer vision, etc...
Fortunately, webcams are becoming more and more prevalent, and there are lots and lots of them all around the world and available for download.

Webcam dataset

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

- 1350 webcams

GPS positions

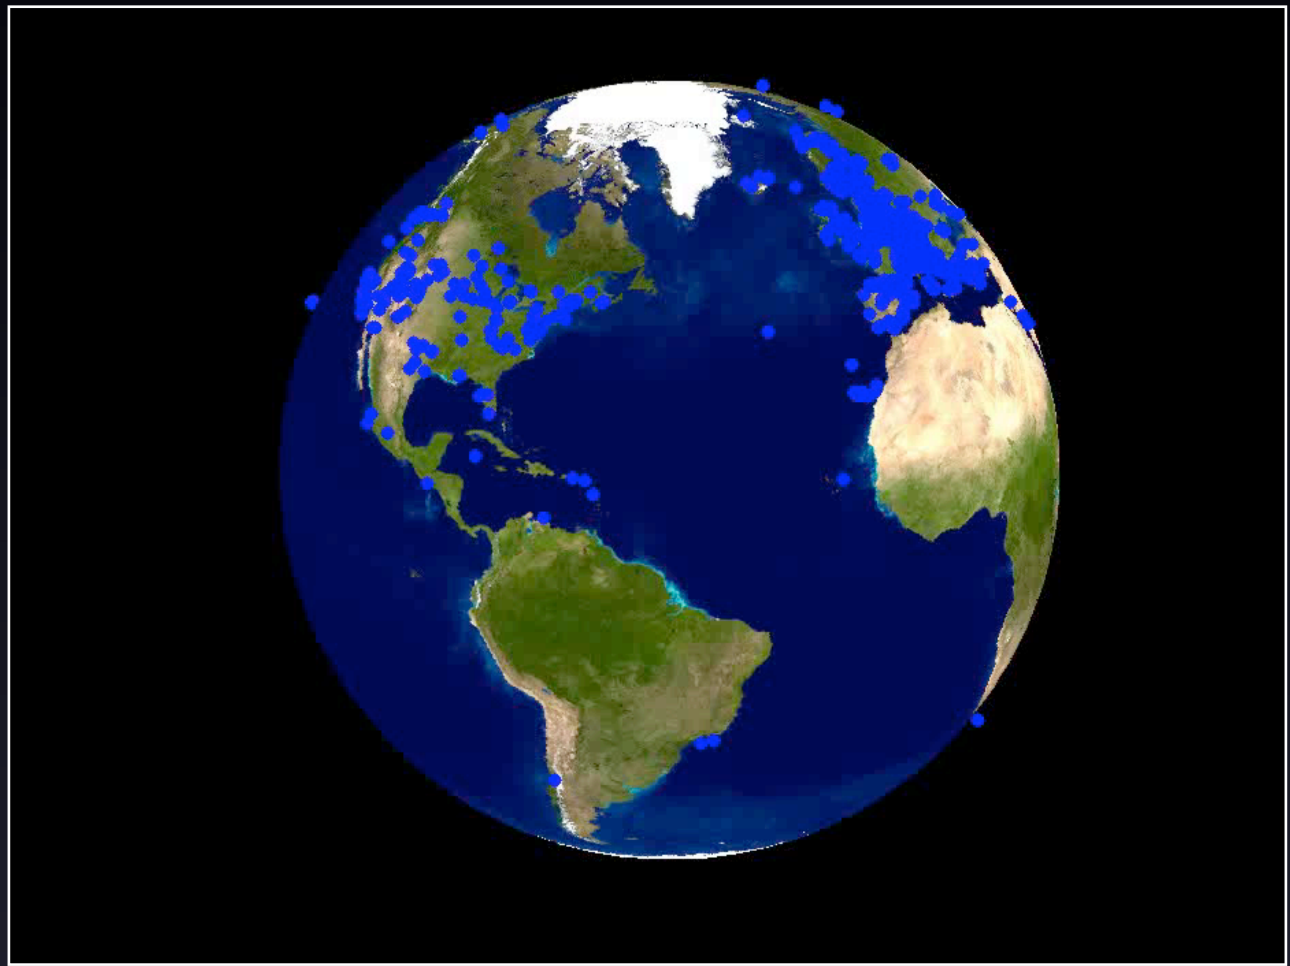

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

- 1350 webcams
- ≥ 640x480 pixels

GPS positions

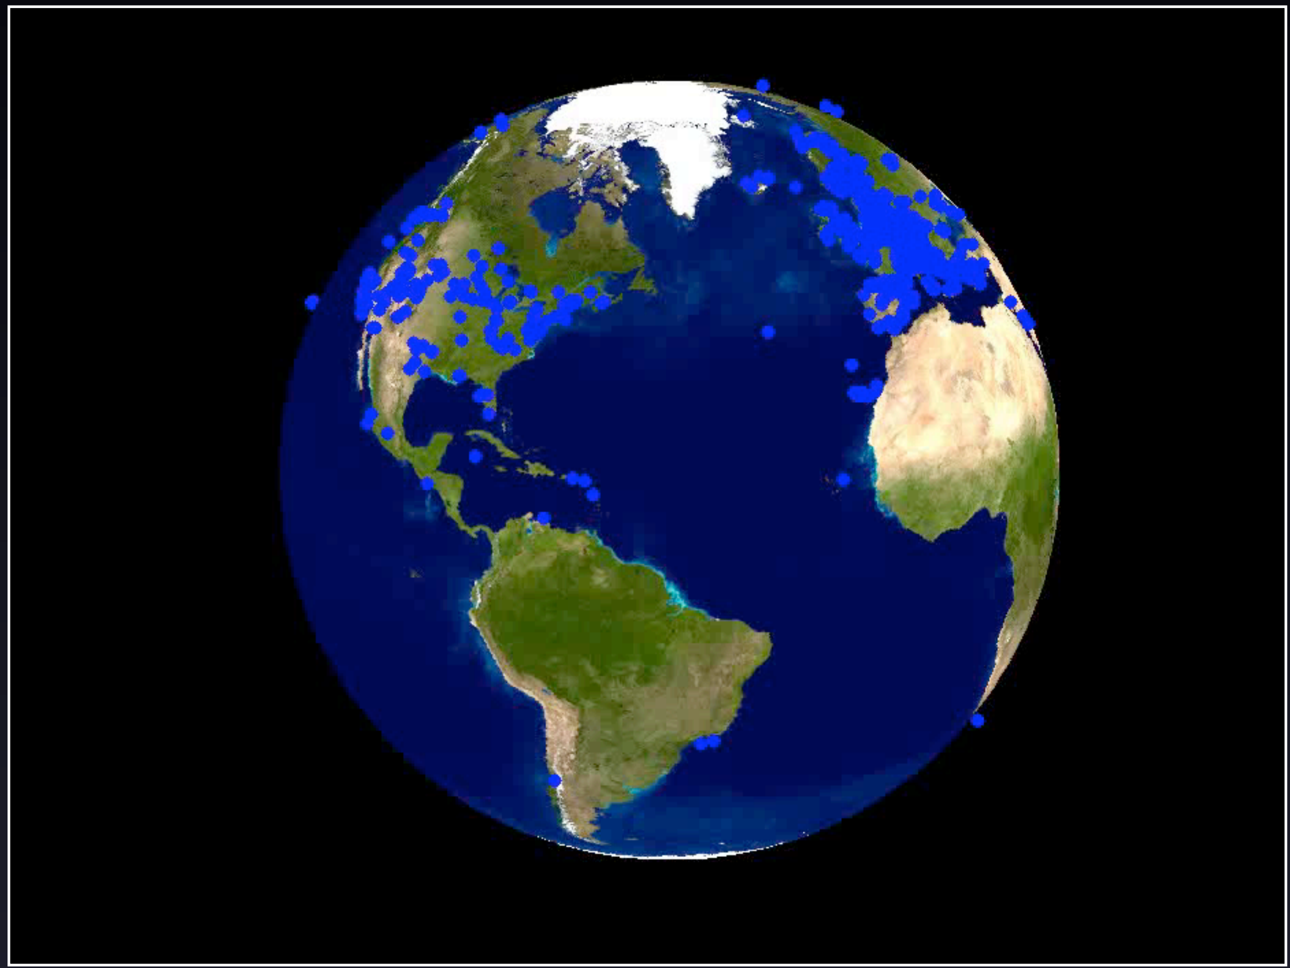

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

- 1350 webcams
- ≥ 640x480 pixels
- every 10 minutes

GPS positions

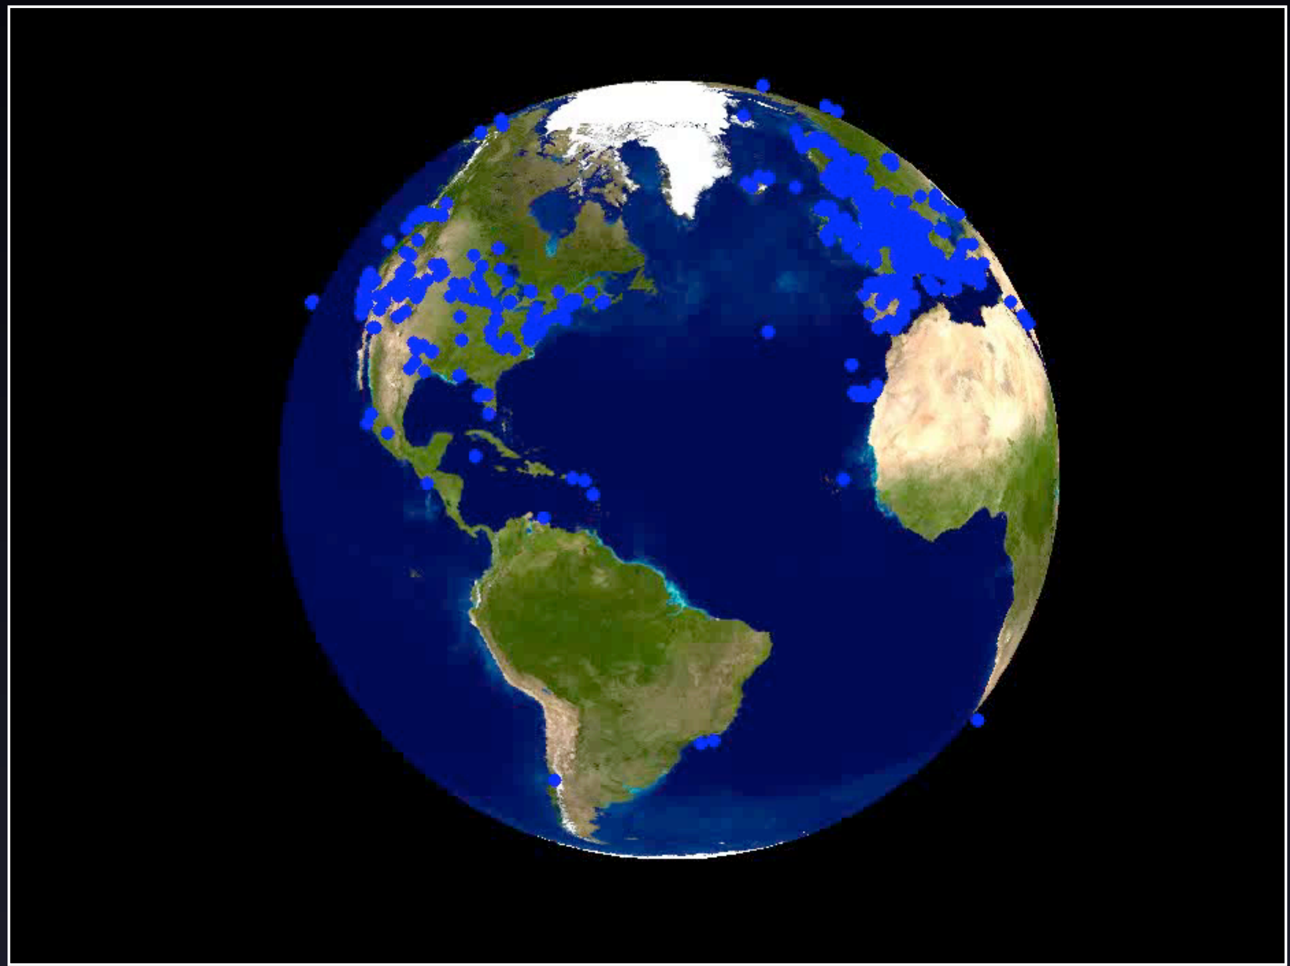

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

- 1350 webcams
- ≥ 640x480 pixels
- every 10 minutes
- 3 months

GPS positions

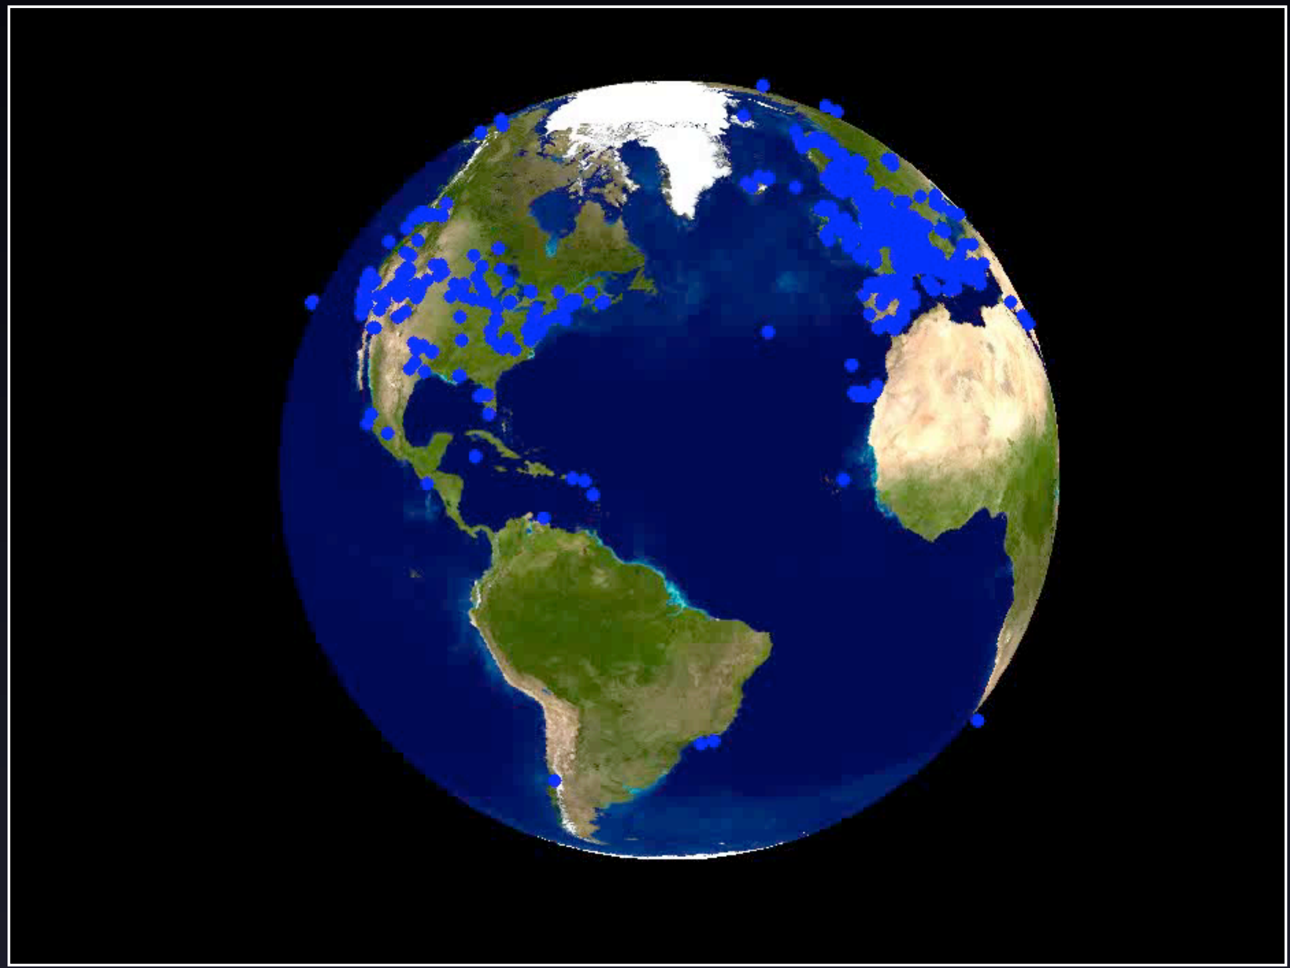

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

- 1350 webcams
- ≥ 640x480 pixels
- every 10 minutes
- 3 months
- 3M images (192 GB)

GPS positions

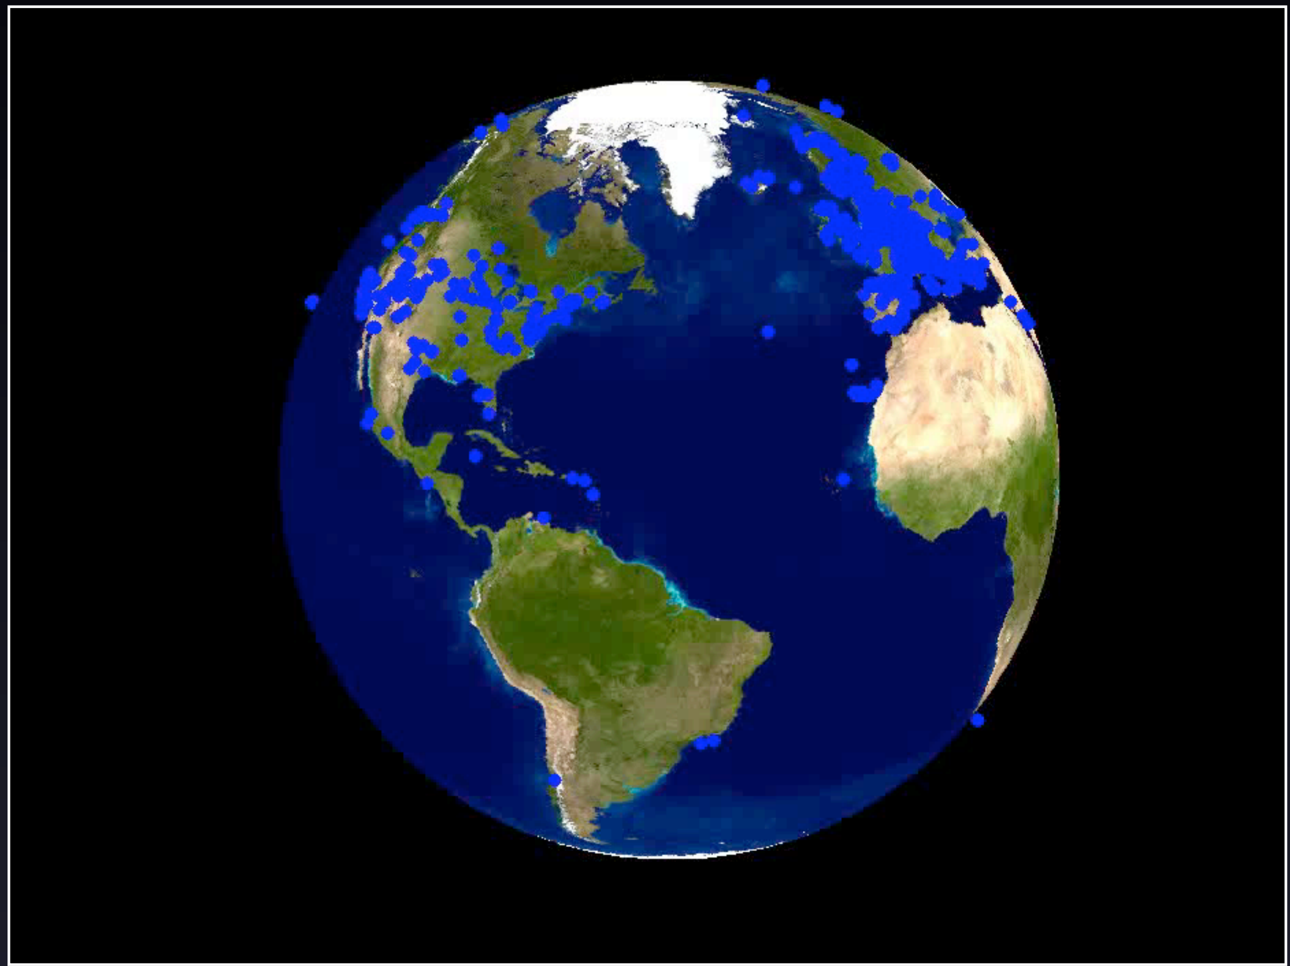

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

Webcam dataset

http://graphics.cs.cmu.edu/projects/webcamdataset

- 1350 webcams
- ≥ 640x480 pixels
- every 10 minutes
- 3 months
- 3M images (192 GB)

GPS positions

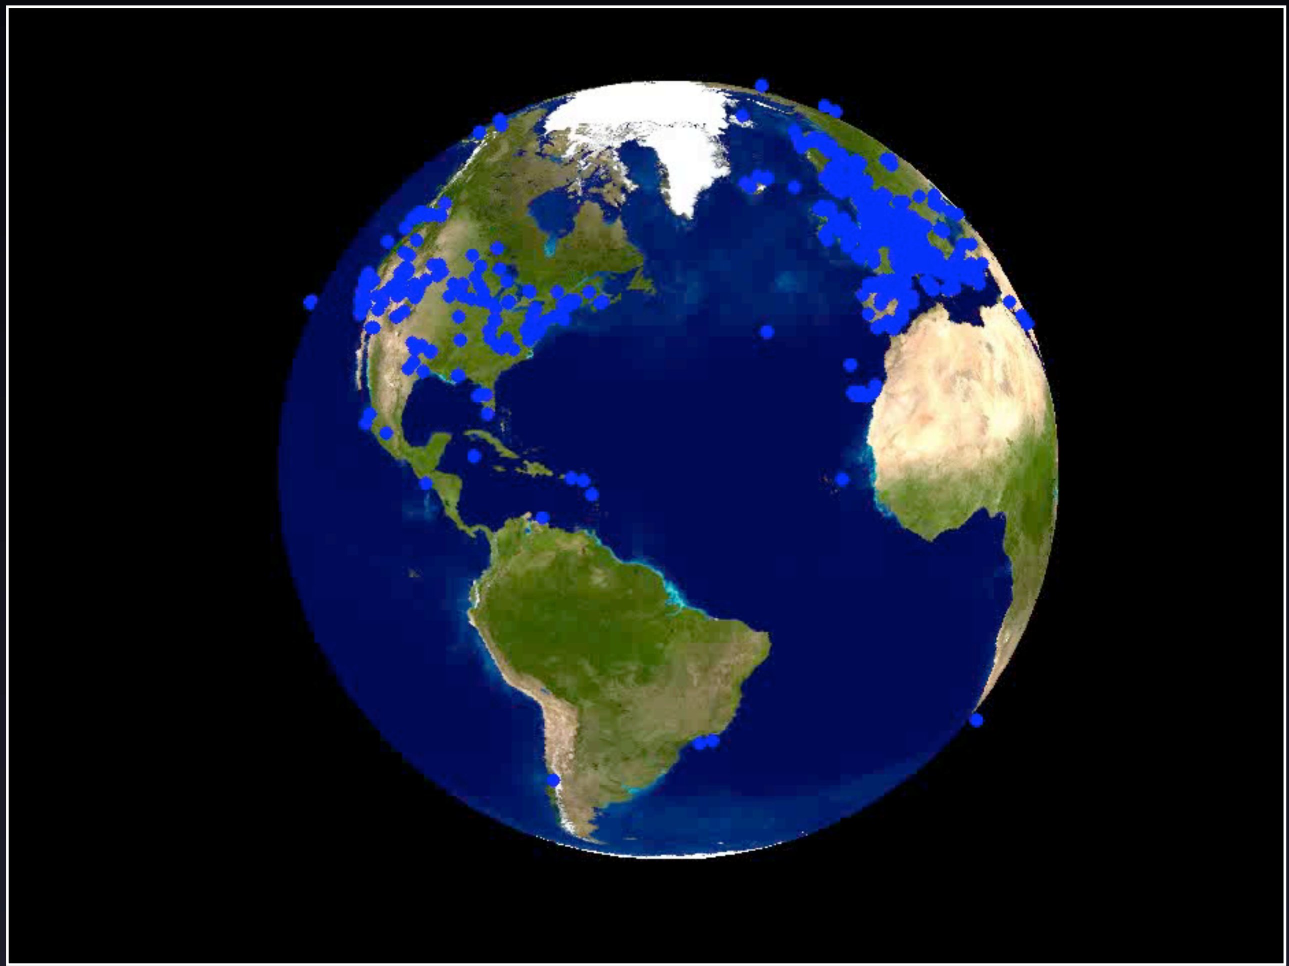

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

We introduce a novel dataset containing more than 1300 high-resolution webcams, located all over the world. All webcams have high enough resolution for graphics, with at least 640x480 pixels. We download images every 10 minutes, and we have roughly 3 months worth of data which amounts to 3M images, but we're going to continue downloading until we get at least 6 months for about 6M images. We're making the dataset freely available for research, and I invite you to take a look at our website at the following URL.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

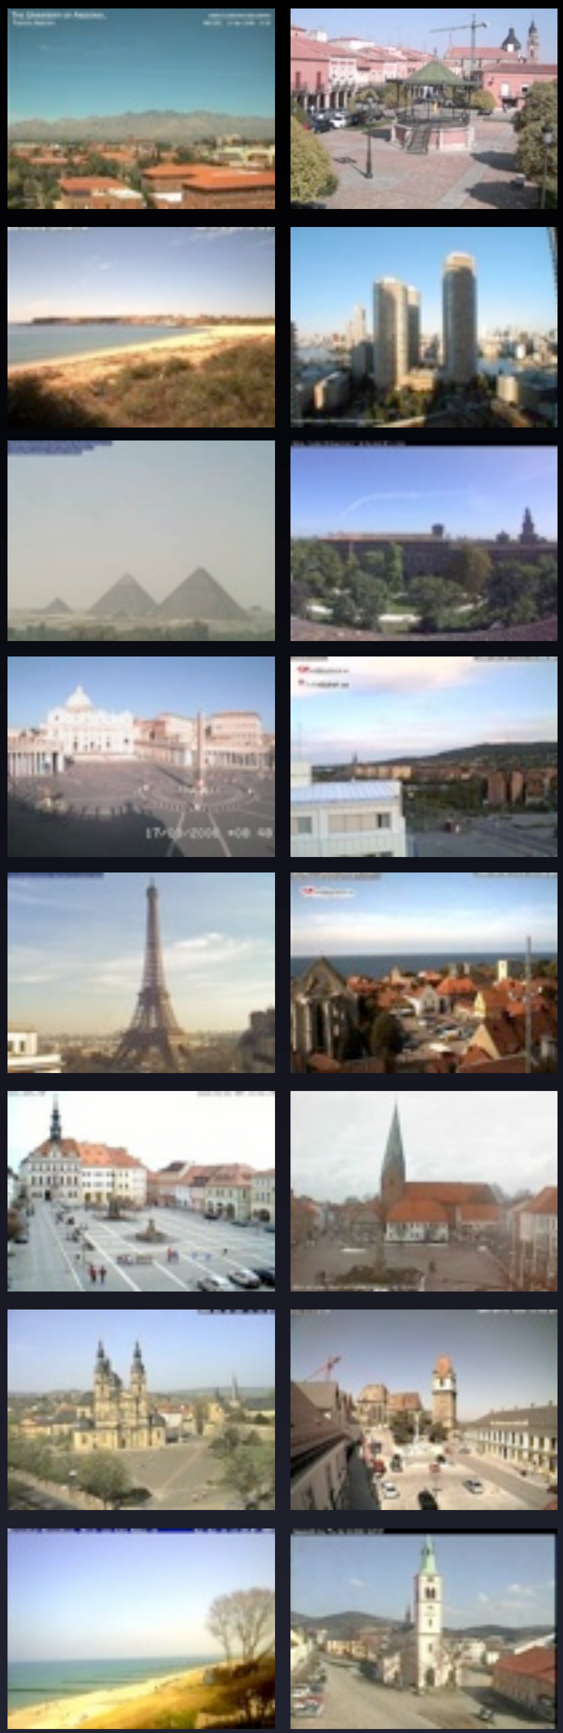

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

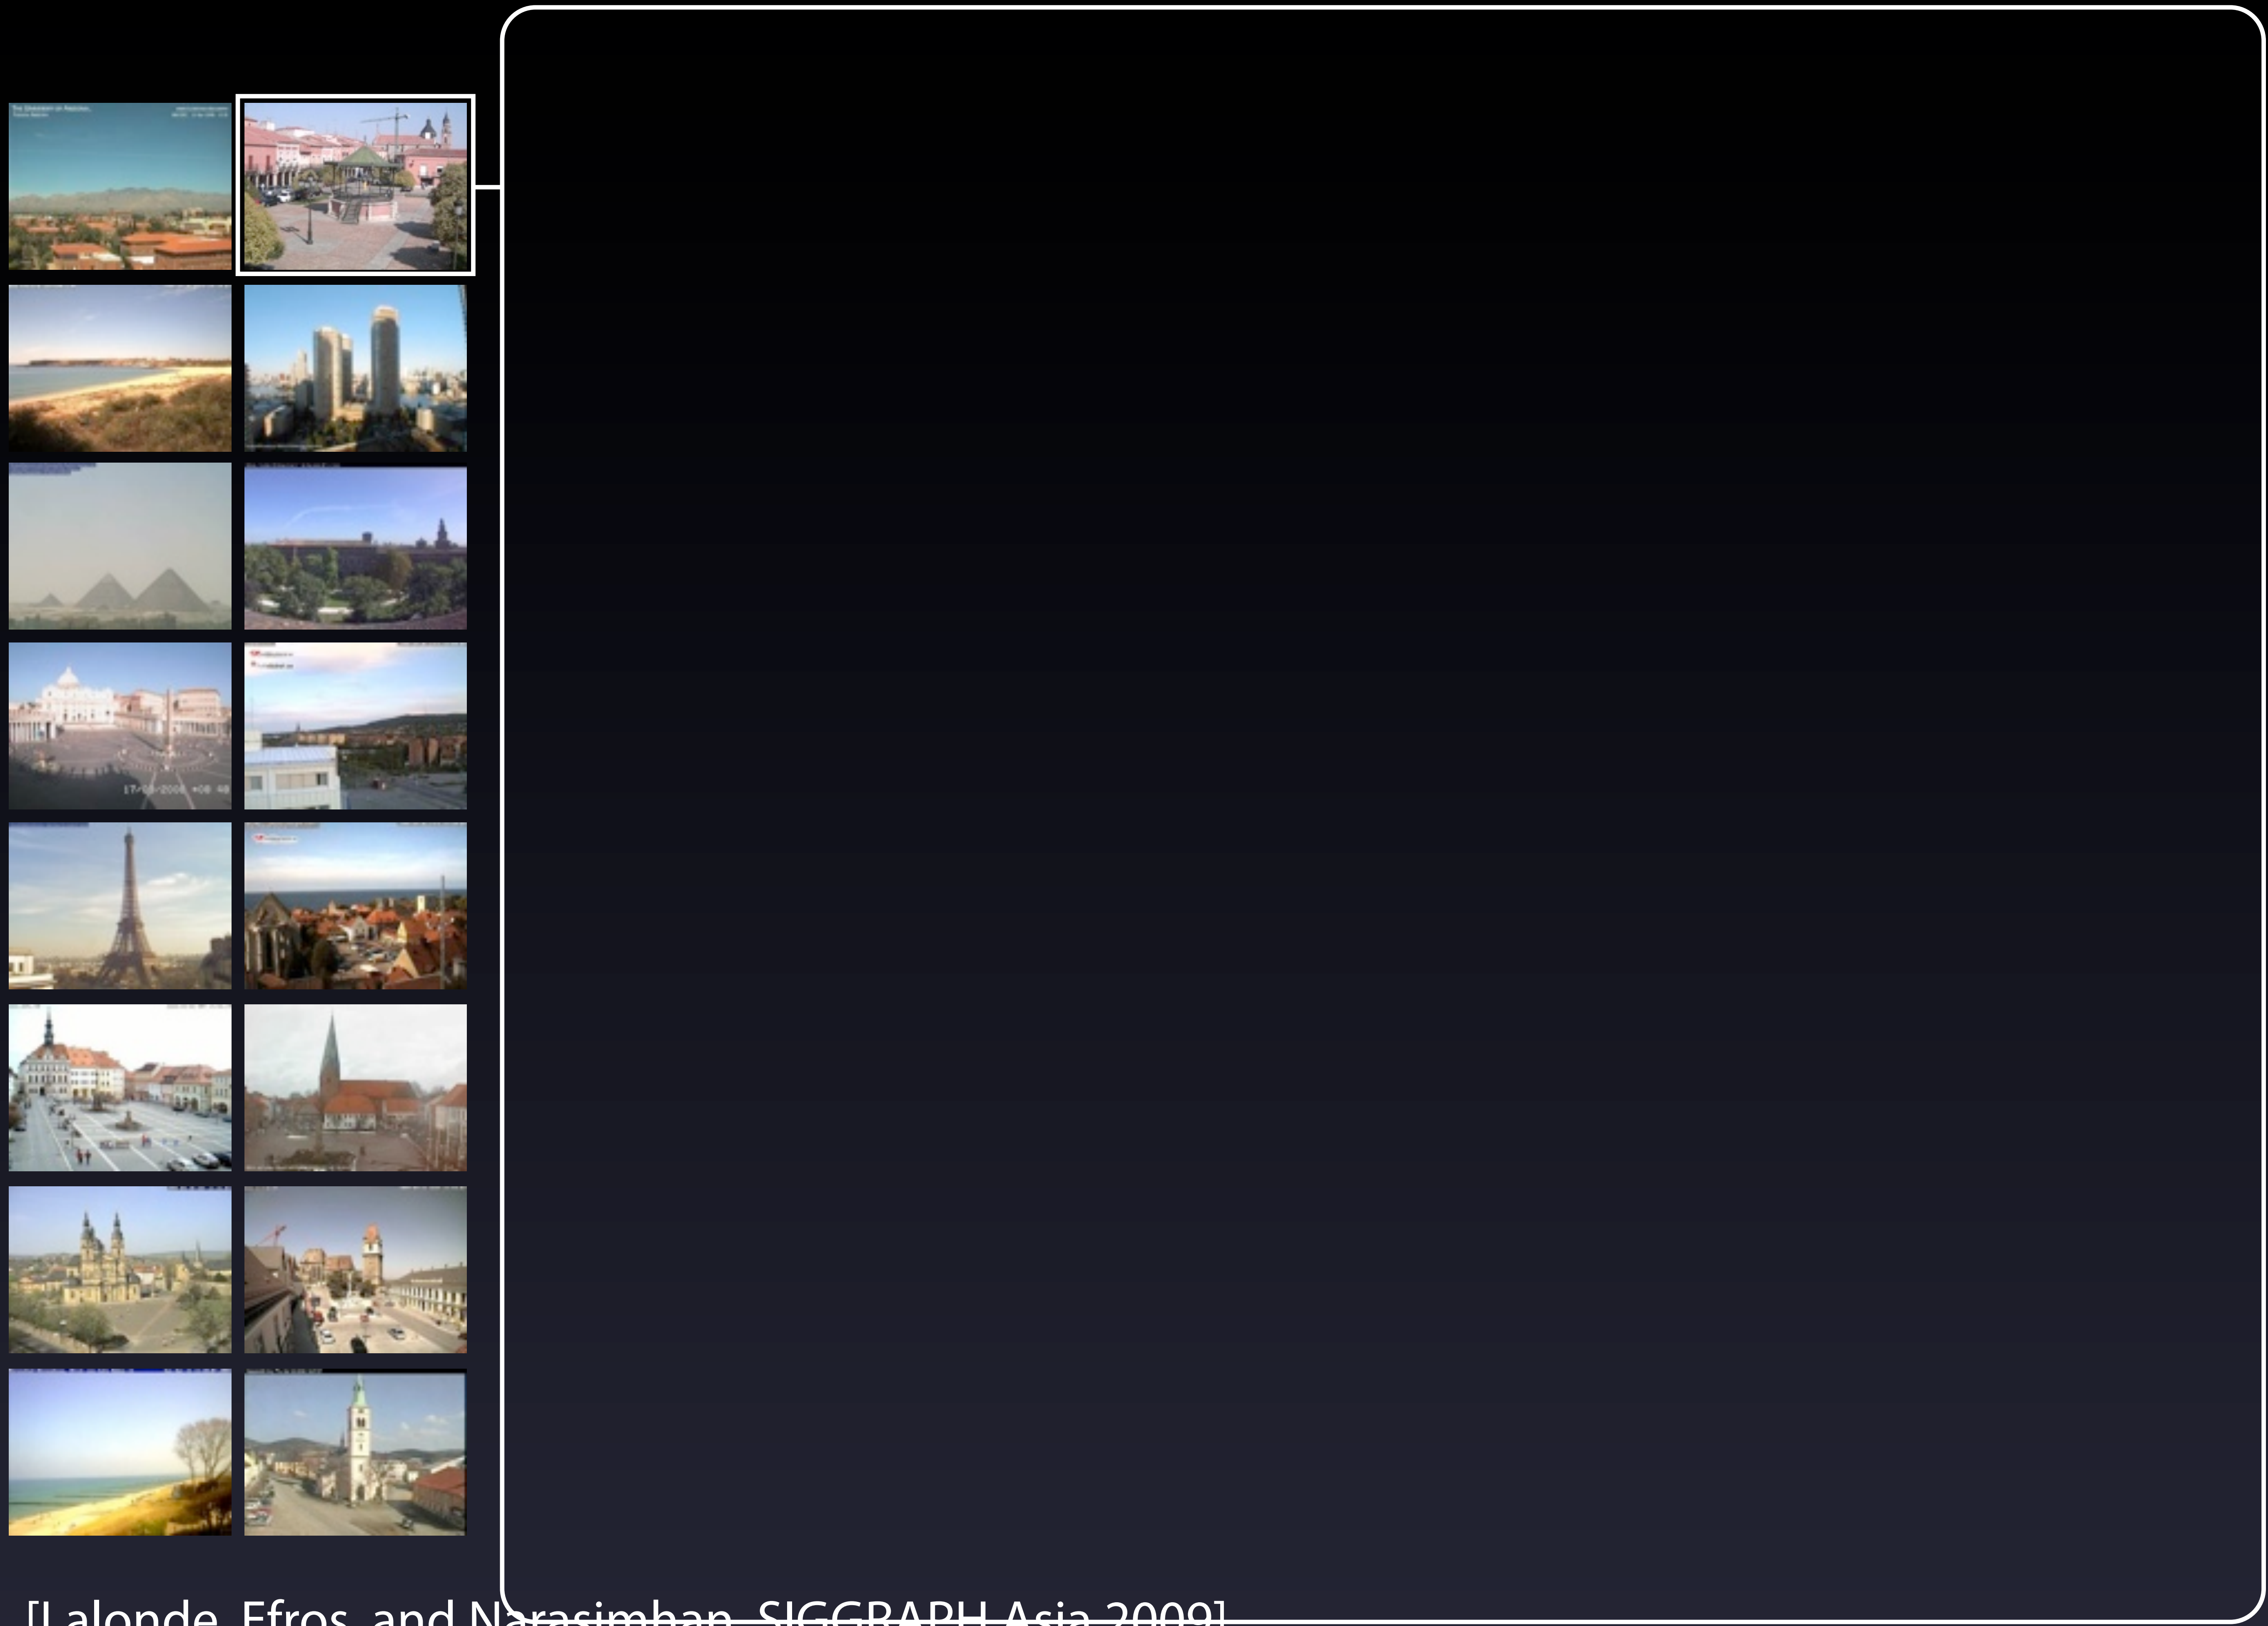

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

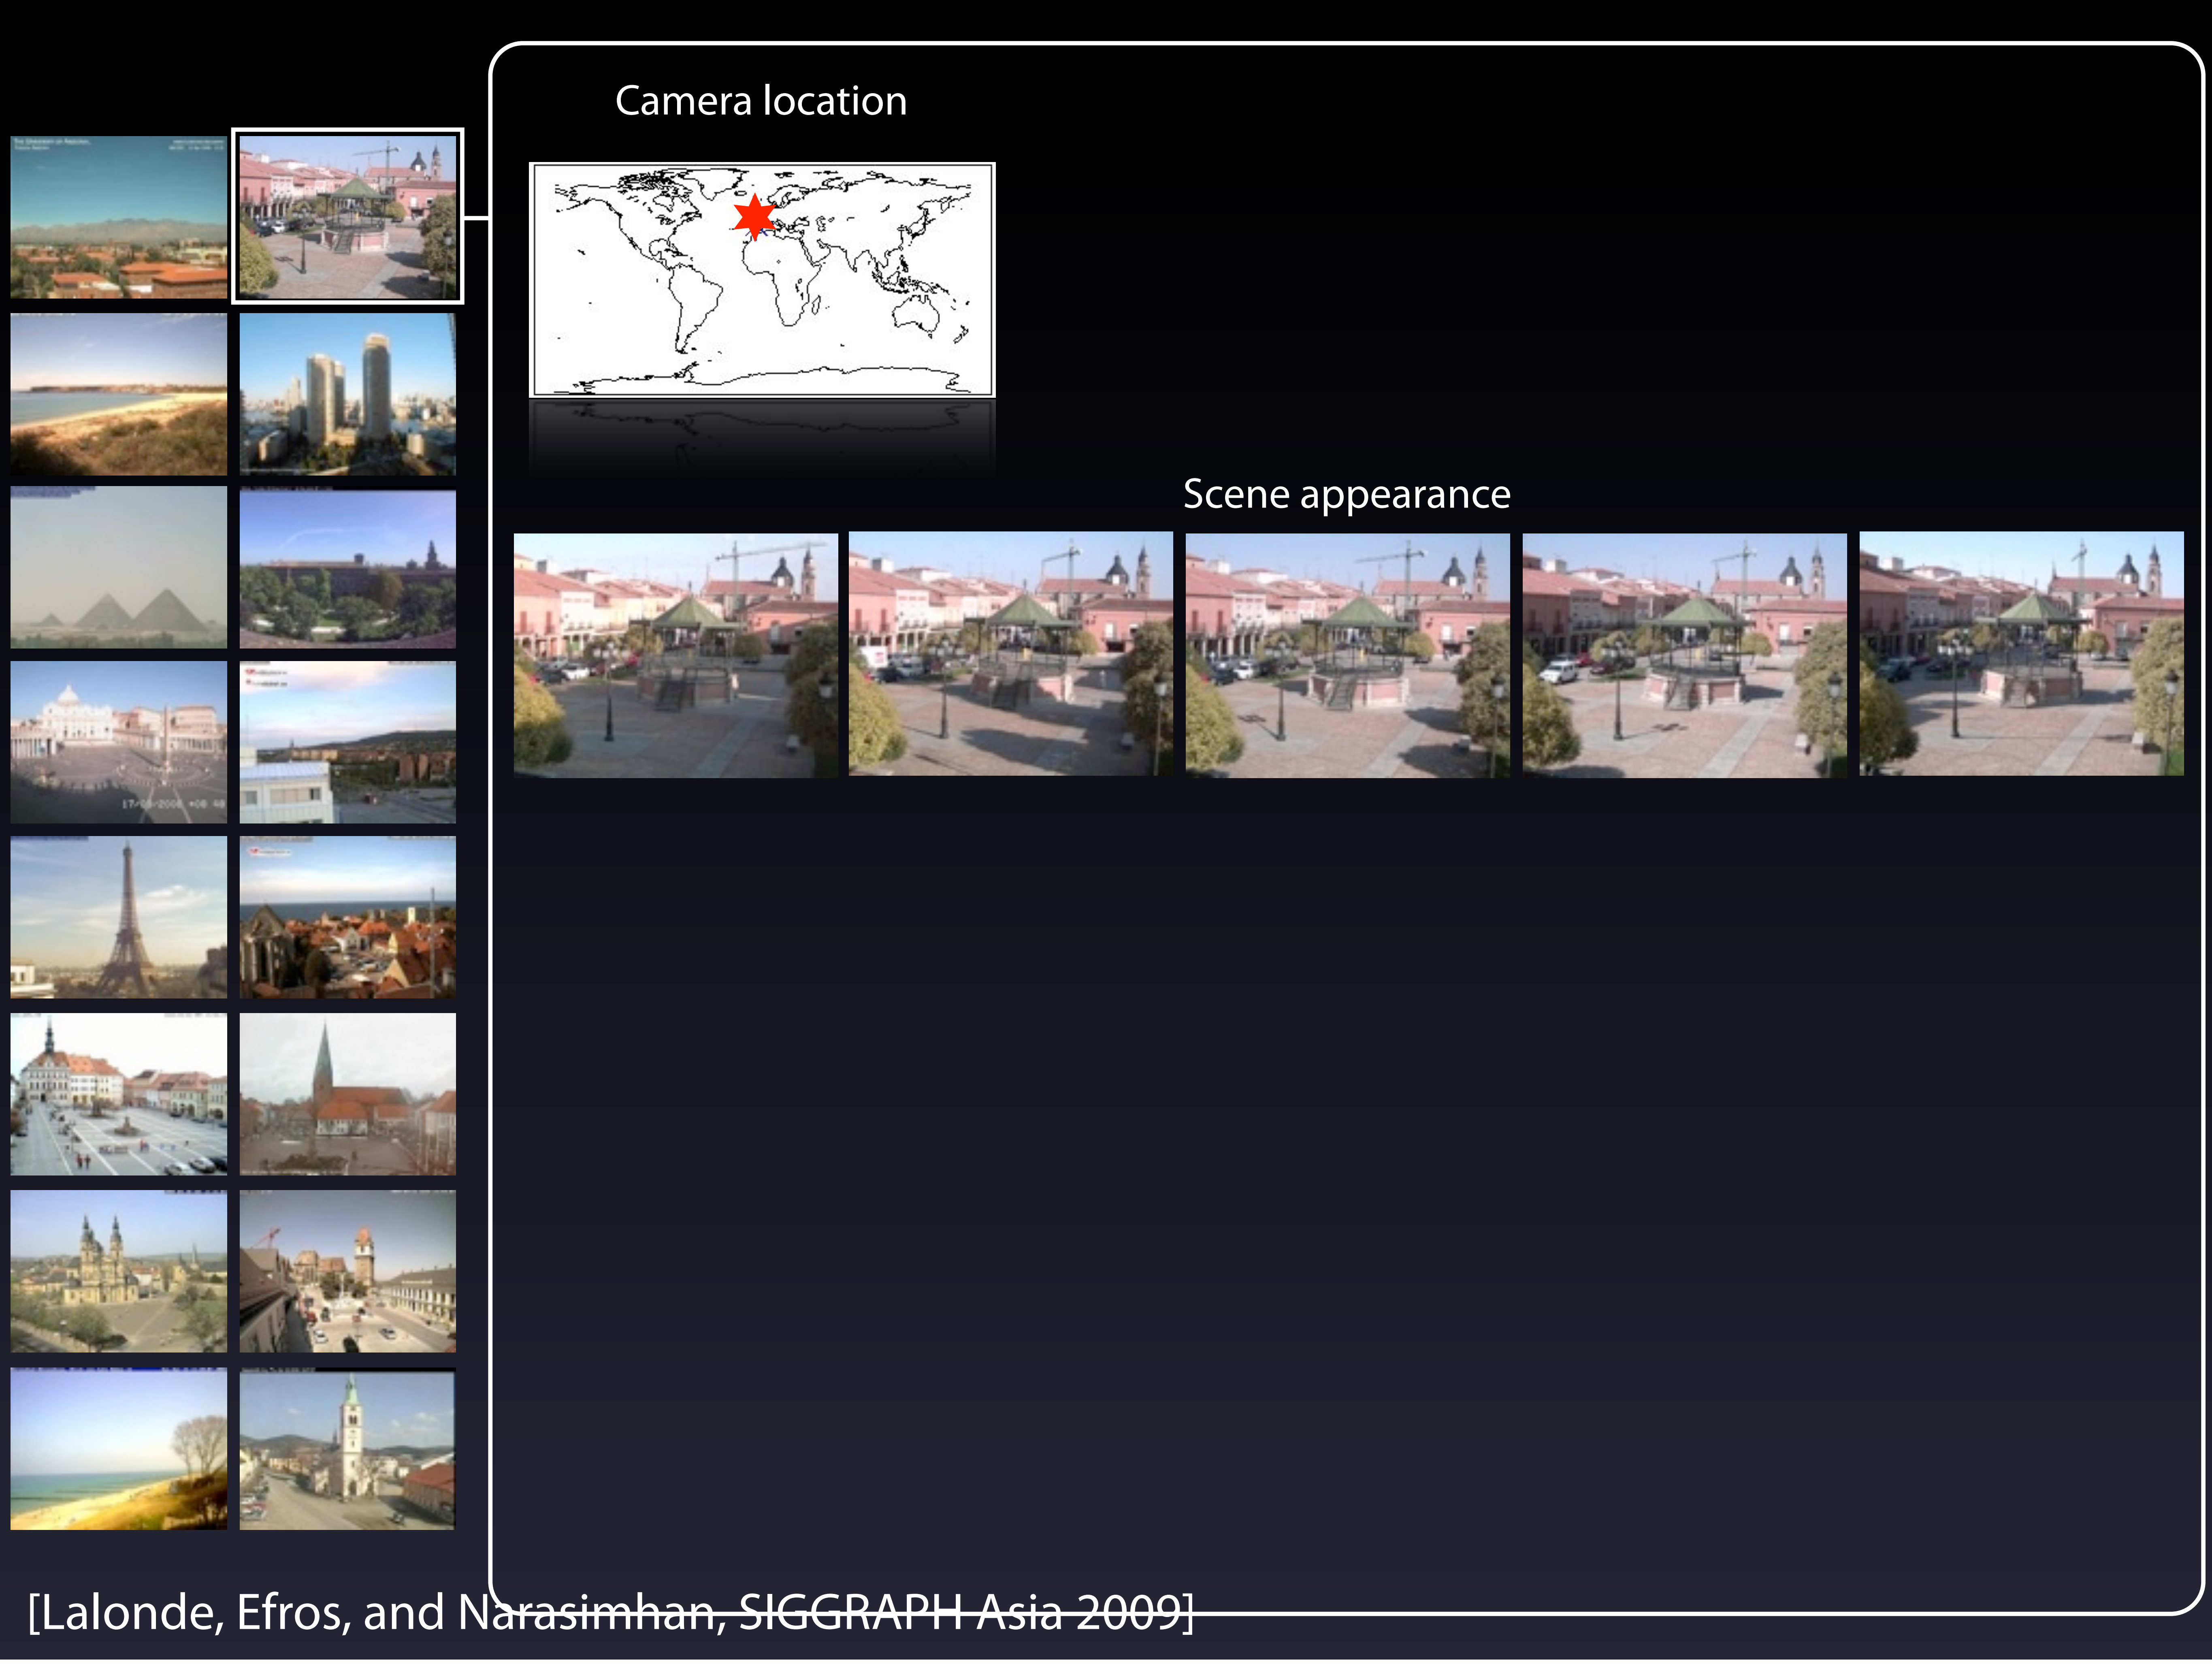

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let’s take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it’s a database of natural scenes and their illumination conditions.

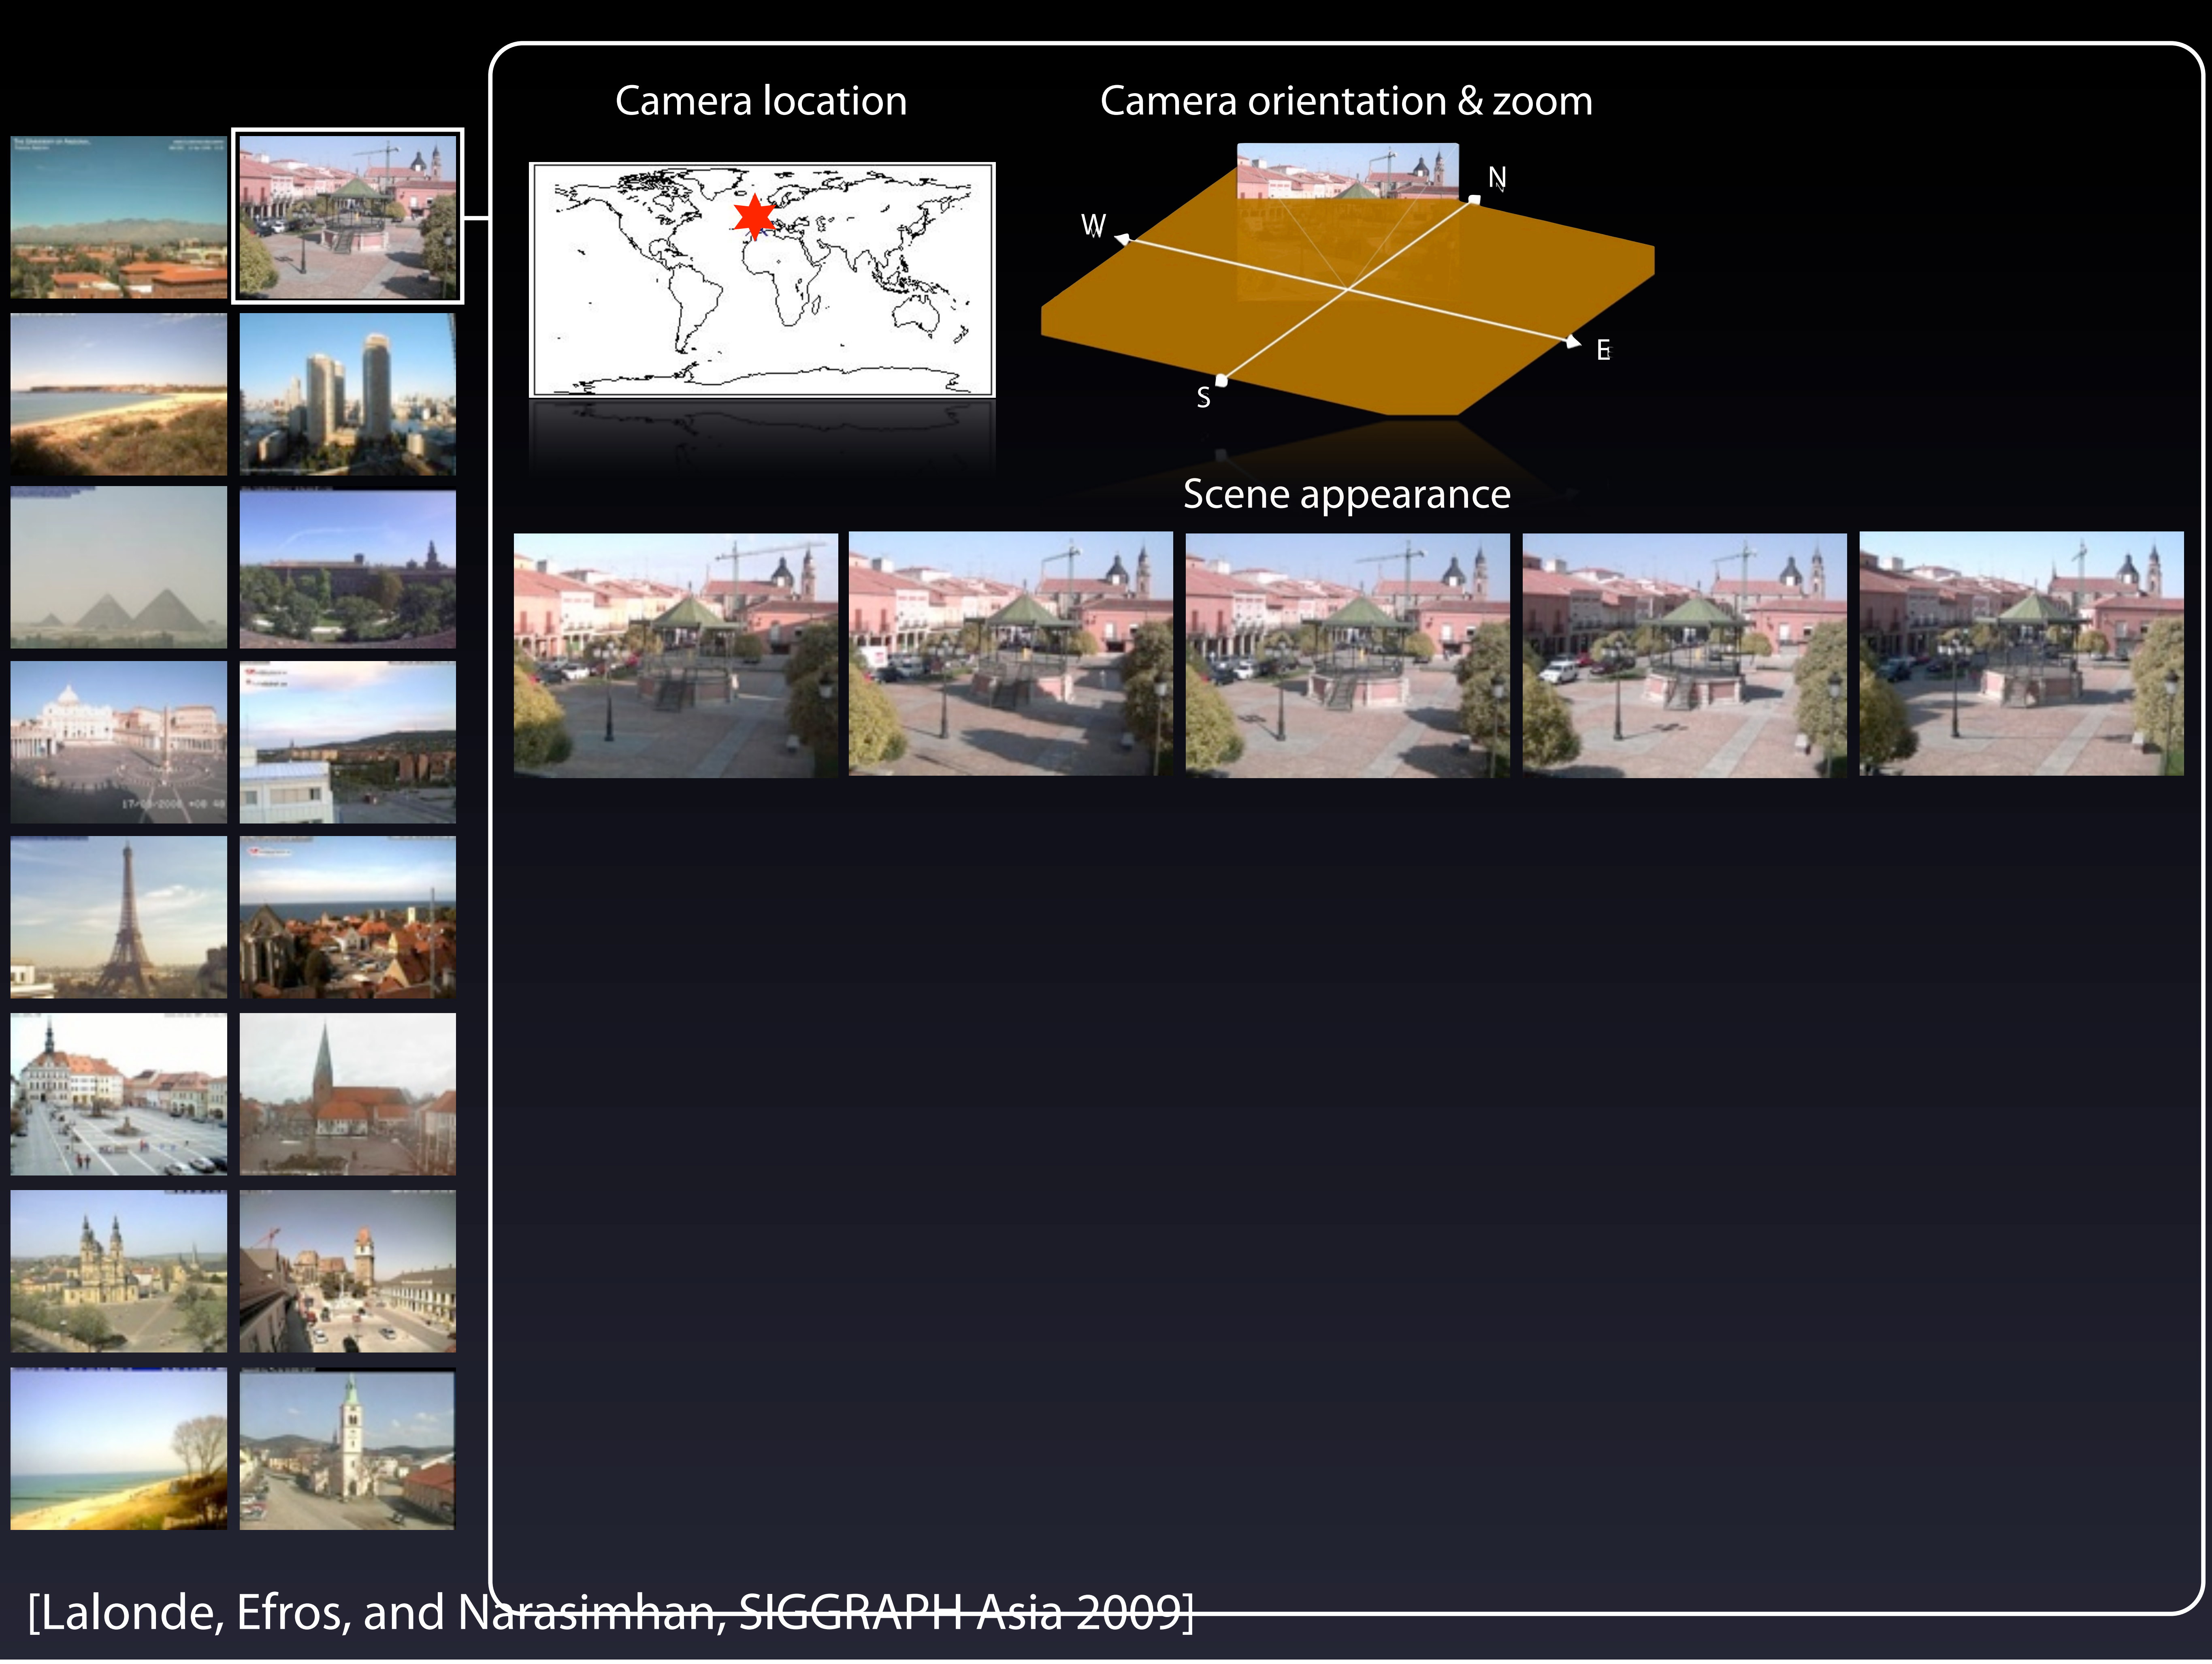

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let’s take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it’s a database of natural scenes and their illumination conditions.

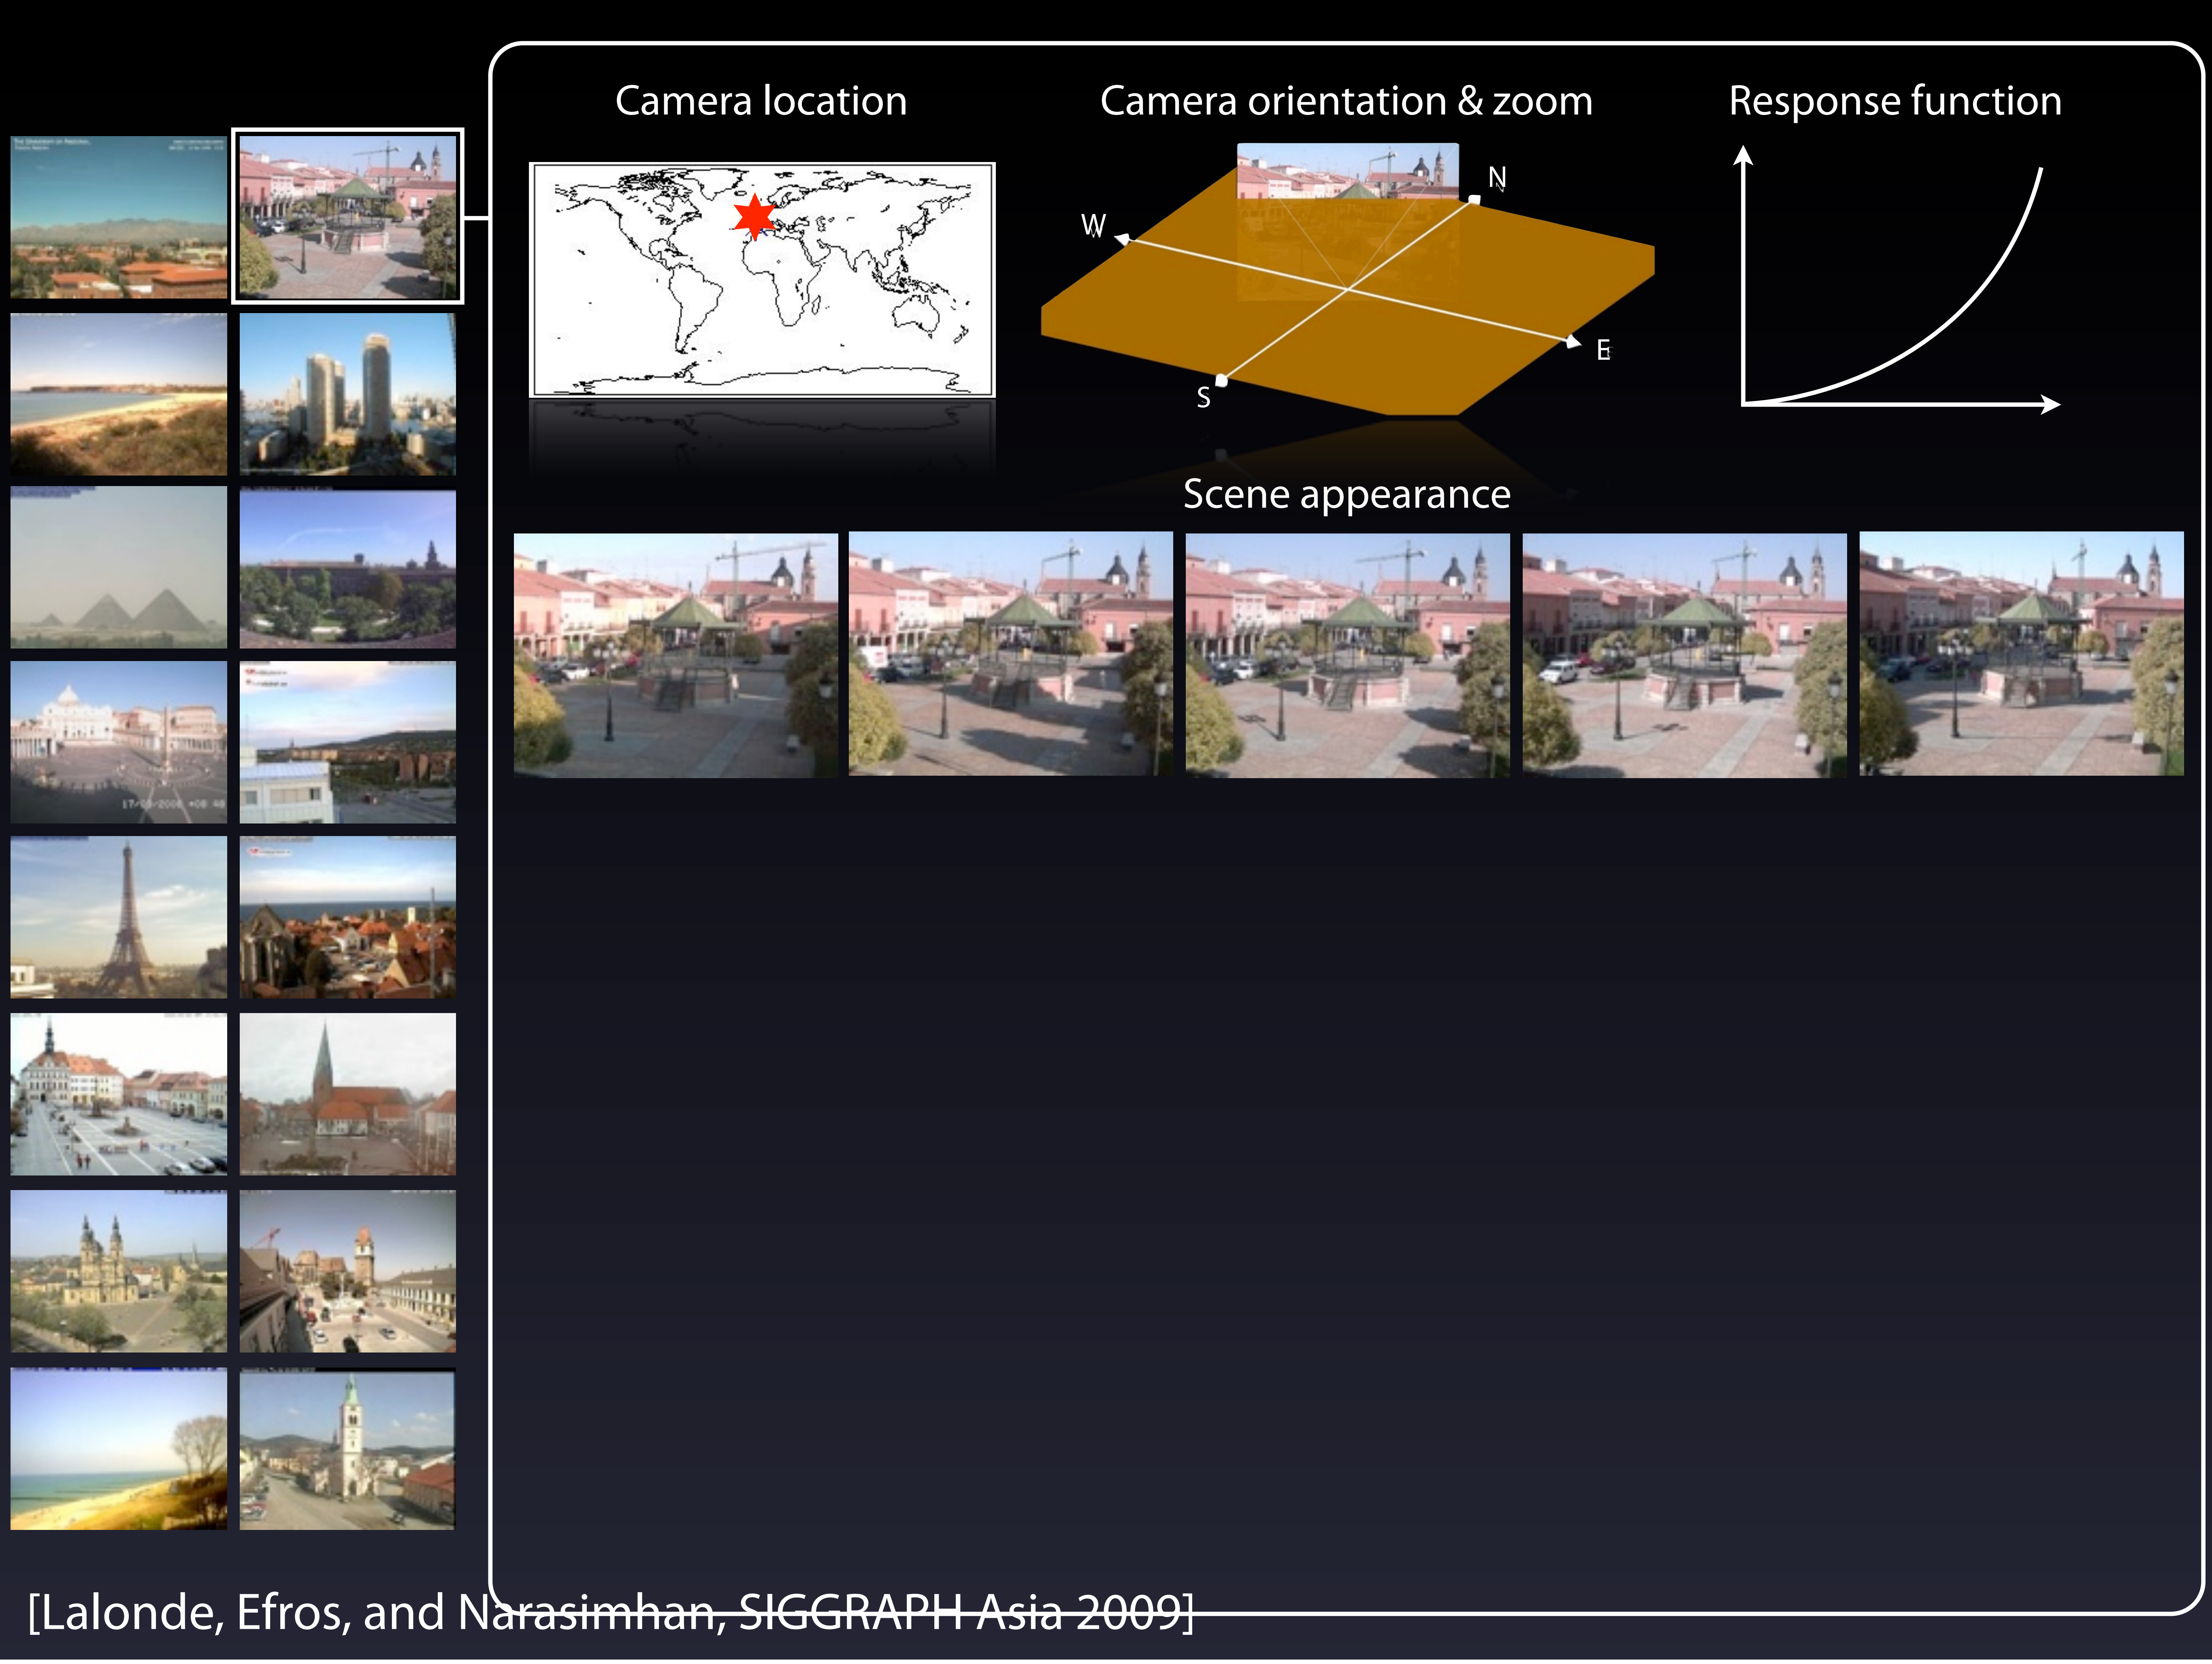

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

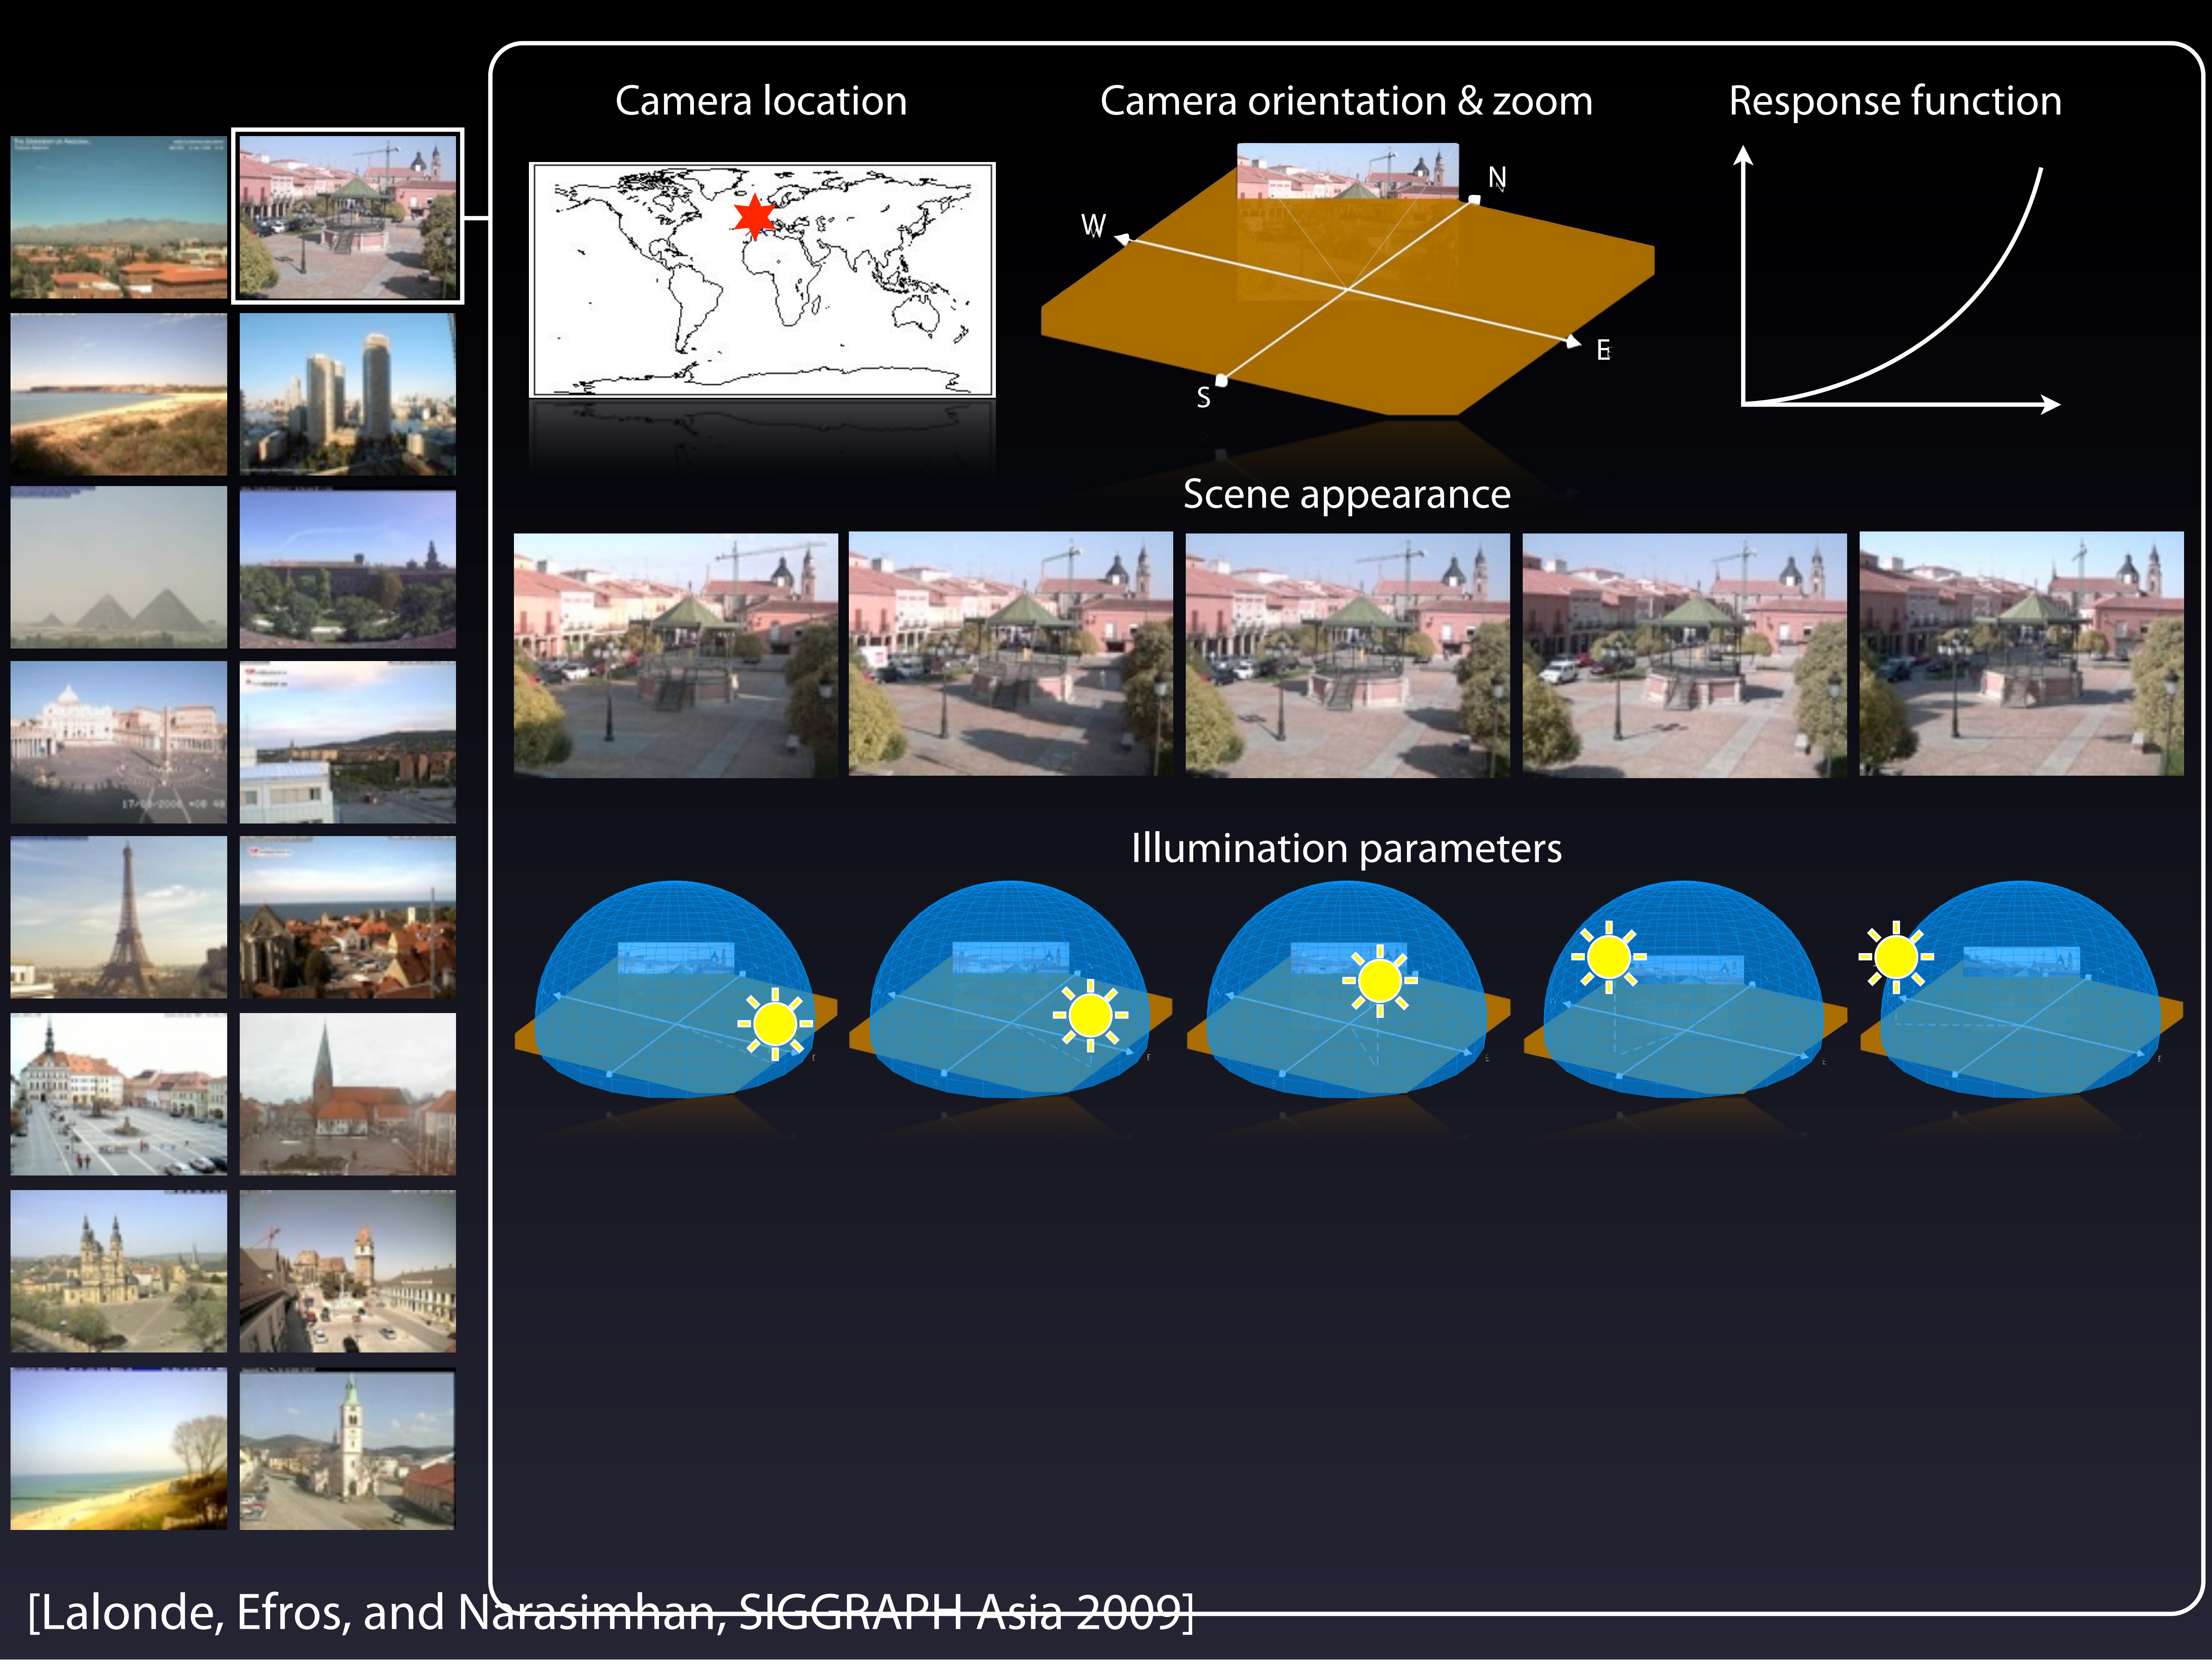

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let’s take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it’s a database of natural scenes and their illumination conditions.

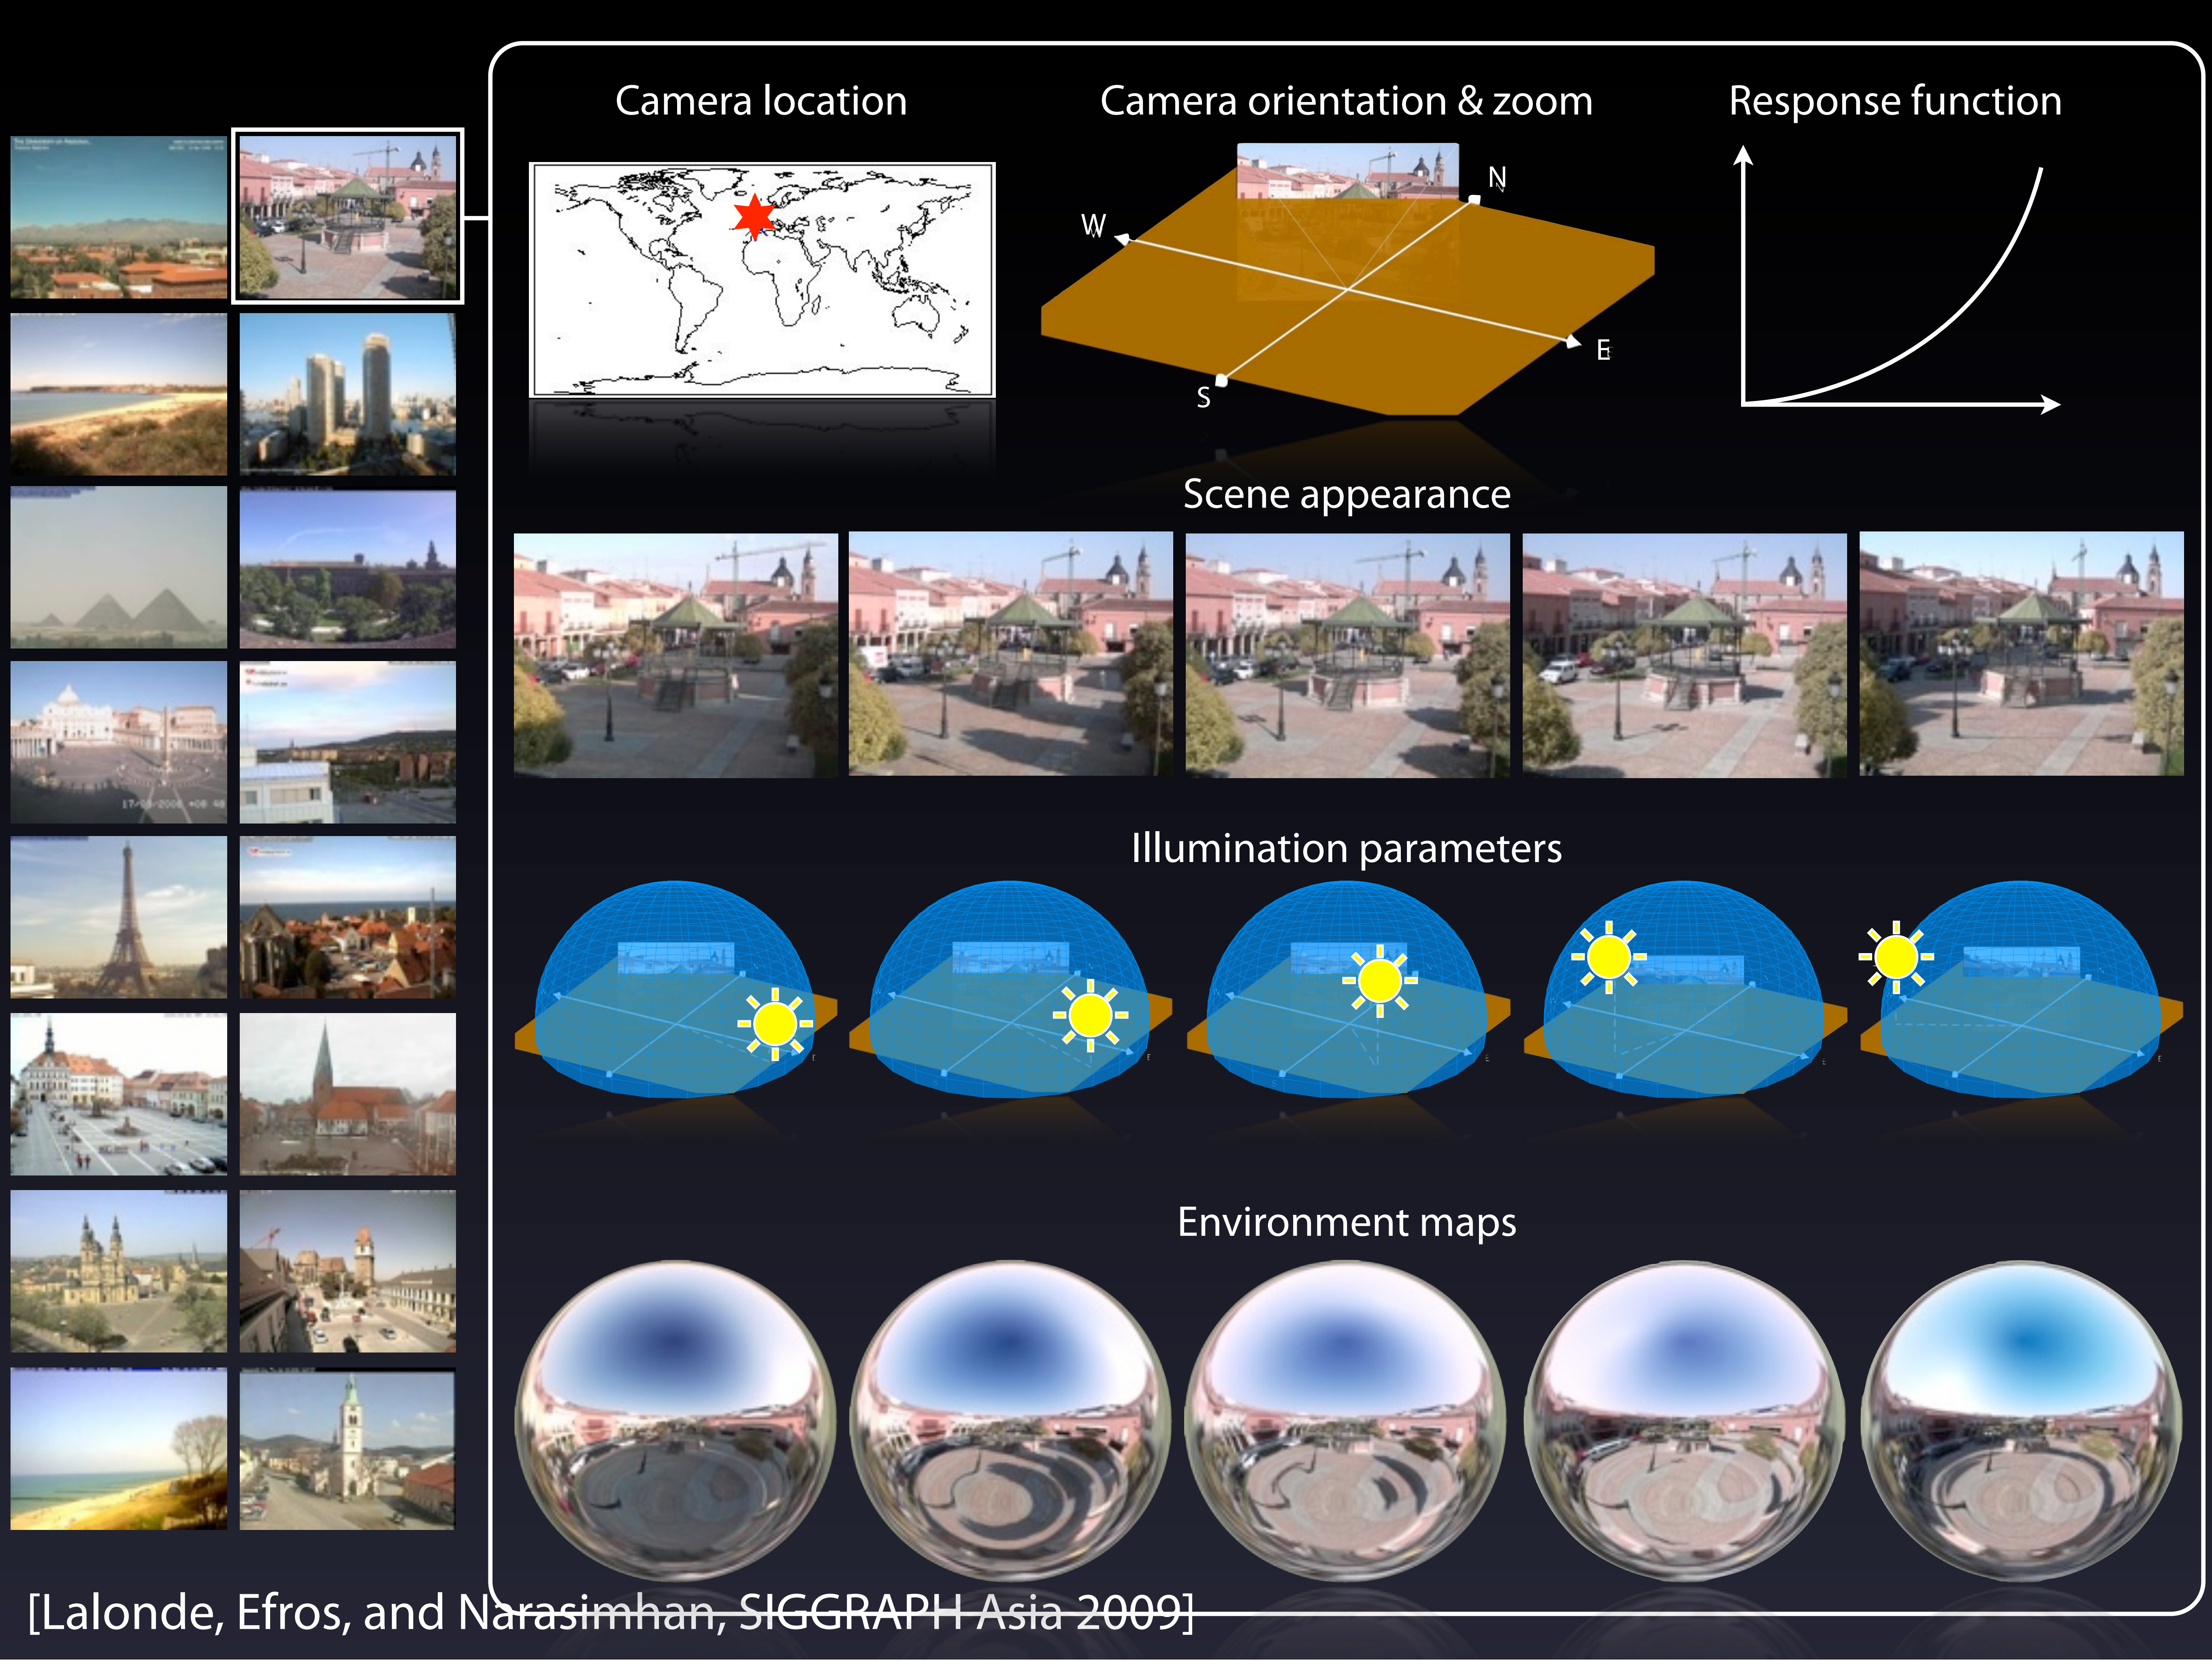

This dataset, combined with the algorithms introduced in this paper, provide very rich information about illumination in natural scenes. Let's take a look at what it contains exactly. When we download the webcams, we end up with thousands of images of the scene under all kinds of illumination conditions. We also get the camera location, and time of capture of each image. We then calibrate the camera and recover its orientation and zoom, and we also estimate the radiometric response function. This allows us to compute the illumination parameters at each image, and extract the corresponding environment maps. As you can see, this is much more than a database of images: it's a database of natural scenes and their illumination conditions.

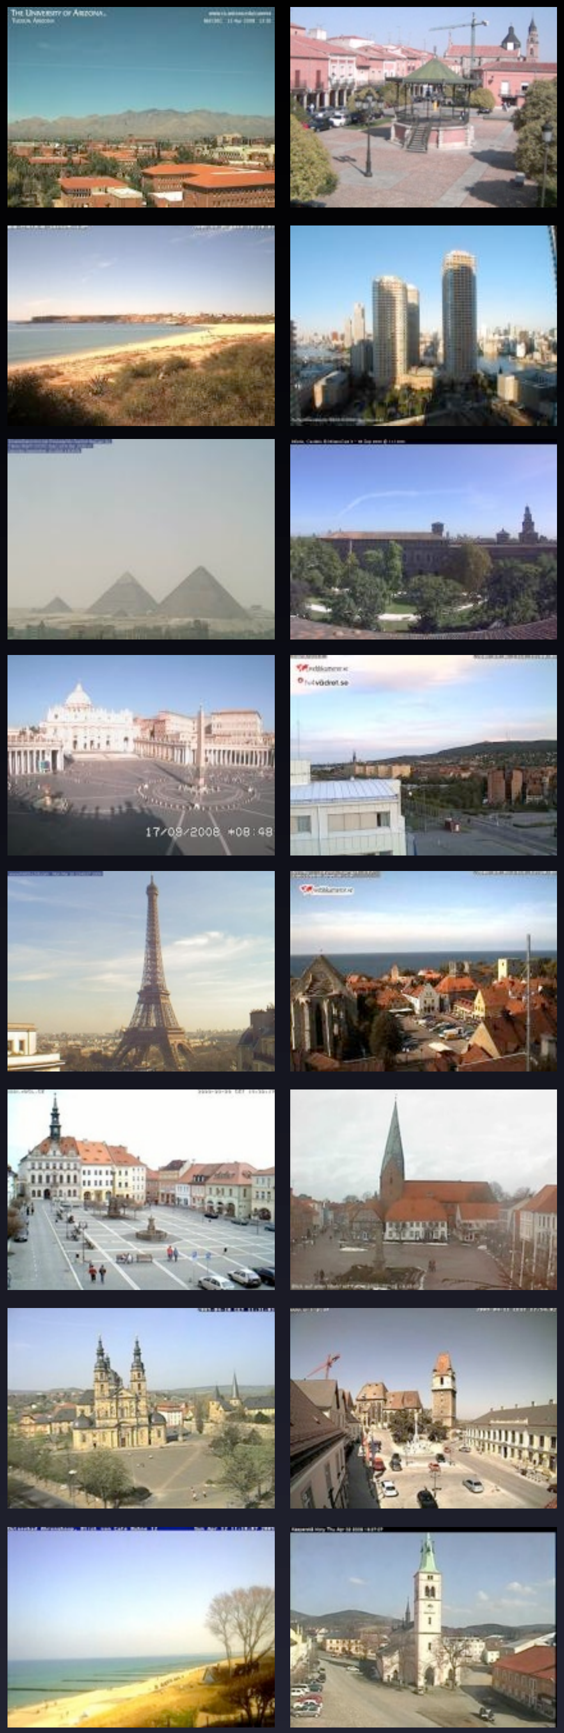

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

but let's take a closer look at the images.

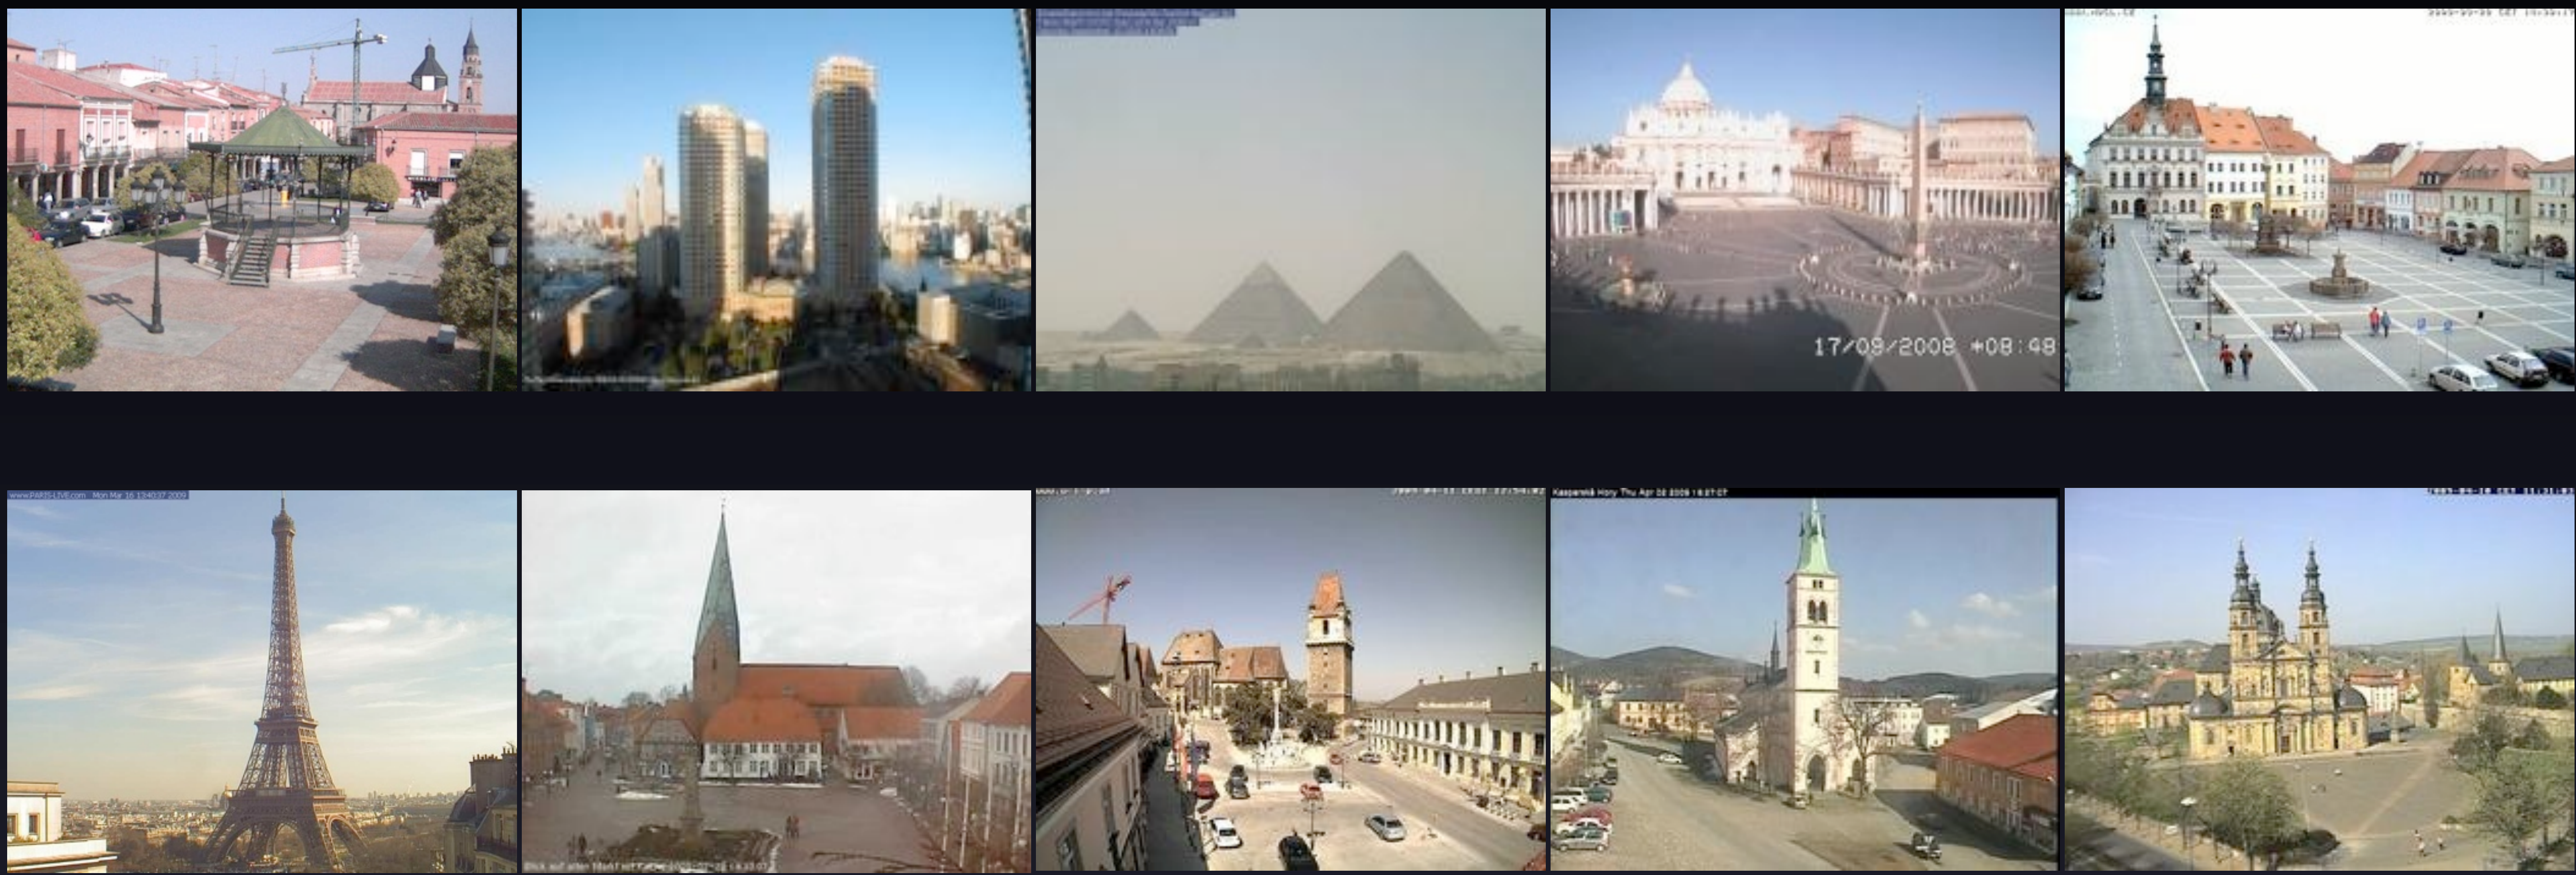

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

but let’s take a closer look at the images.

Illumination-aware objects

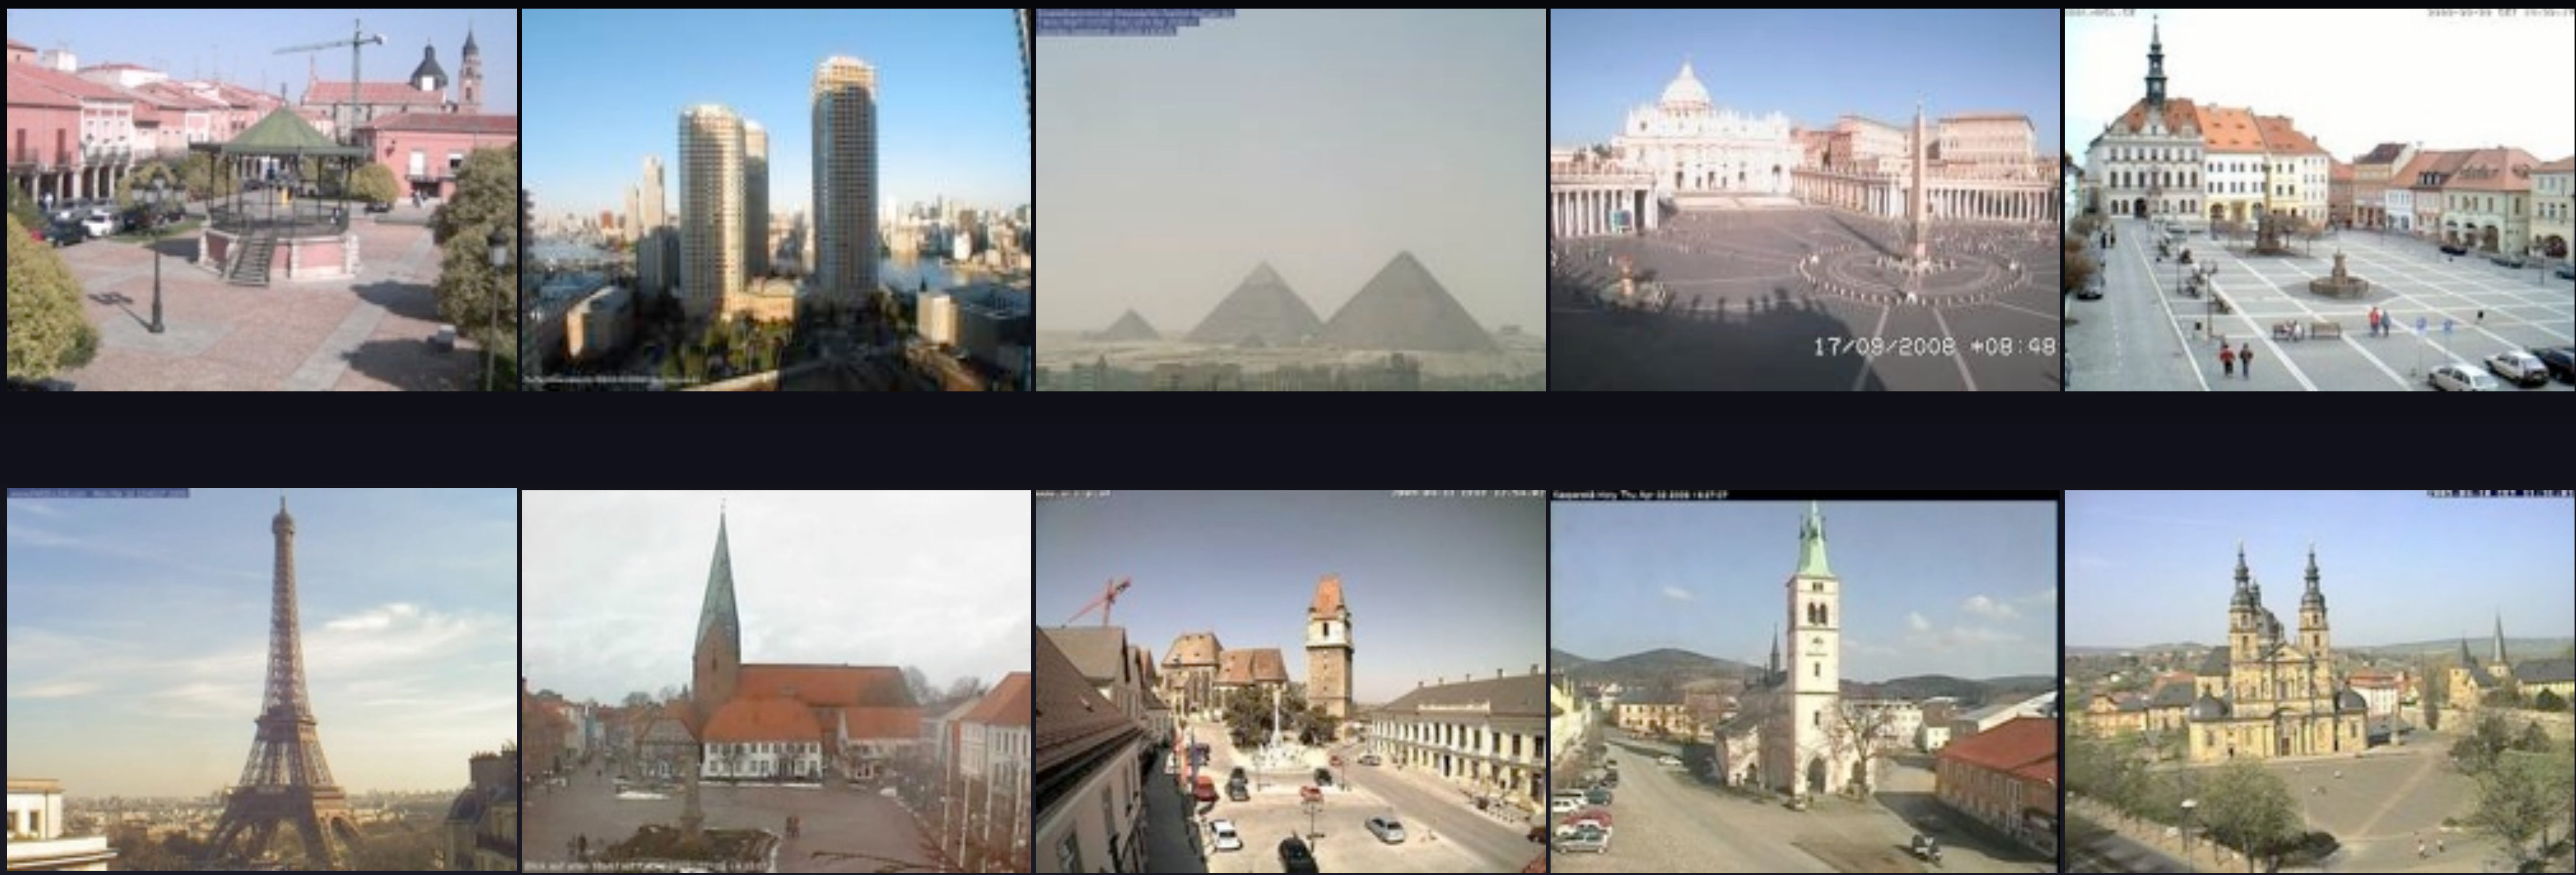

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Most webcams are not just empty scenes, but also contain interesting objects: architectural landmarks, buildings, etc... So it’s also more than a database of scenes, but a database of objects, for which we have

Illumination-aware objects

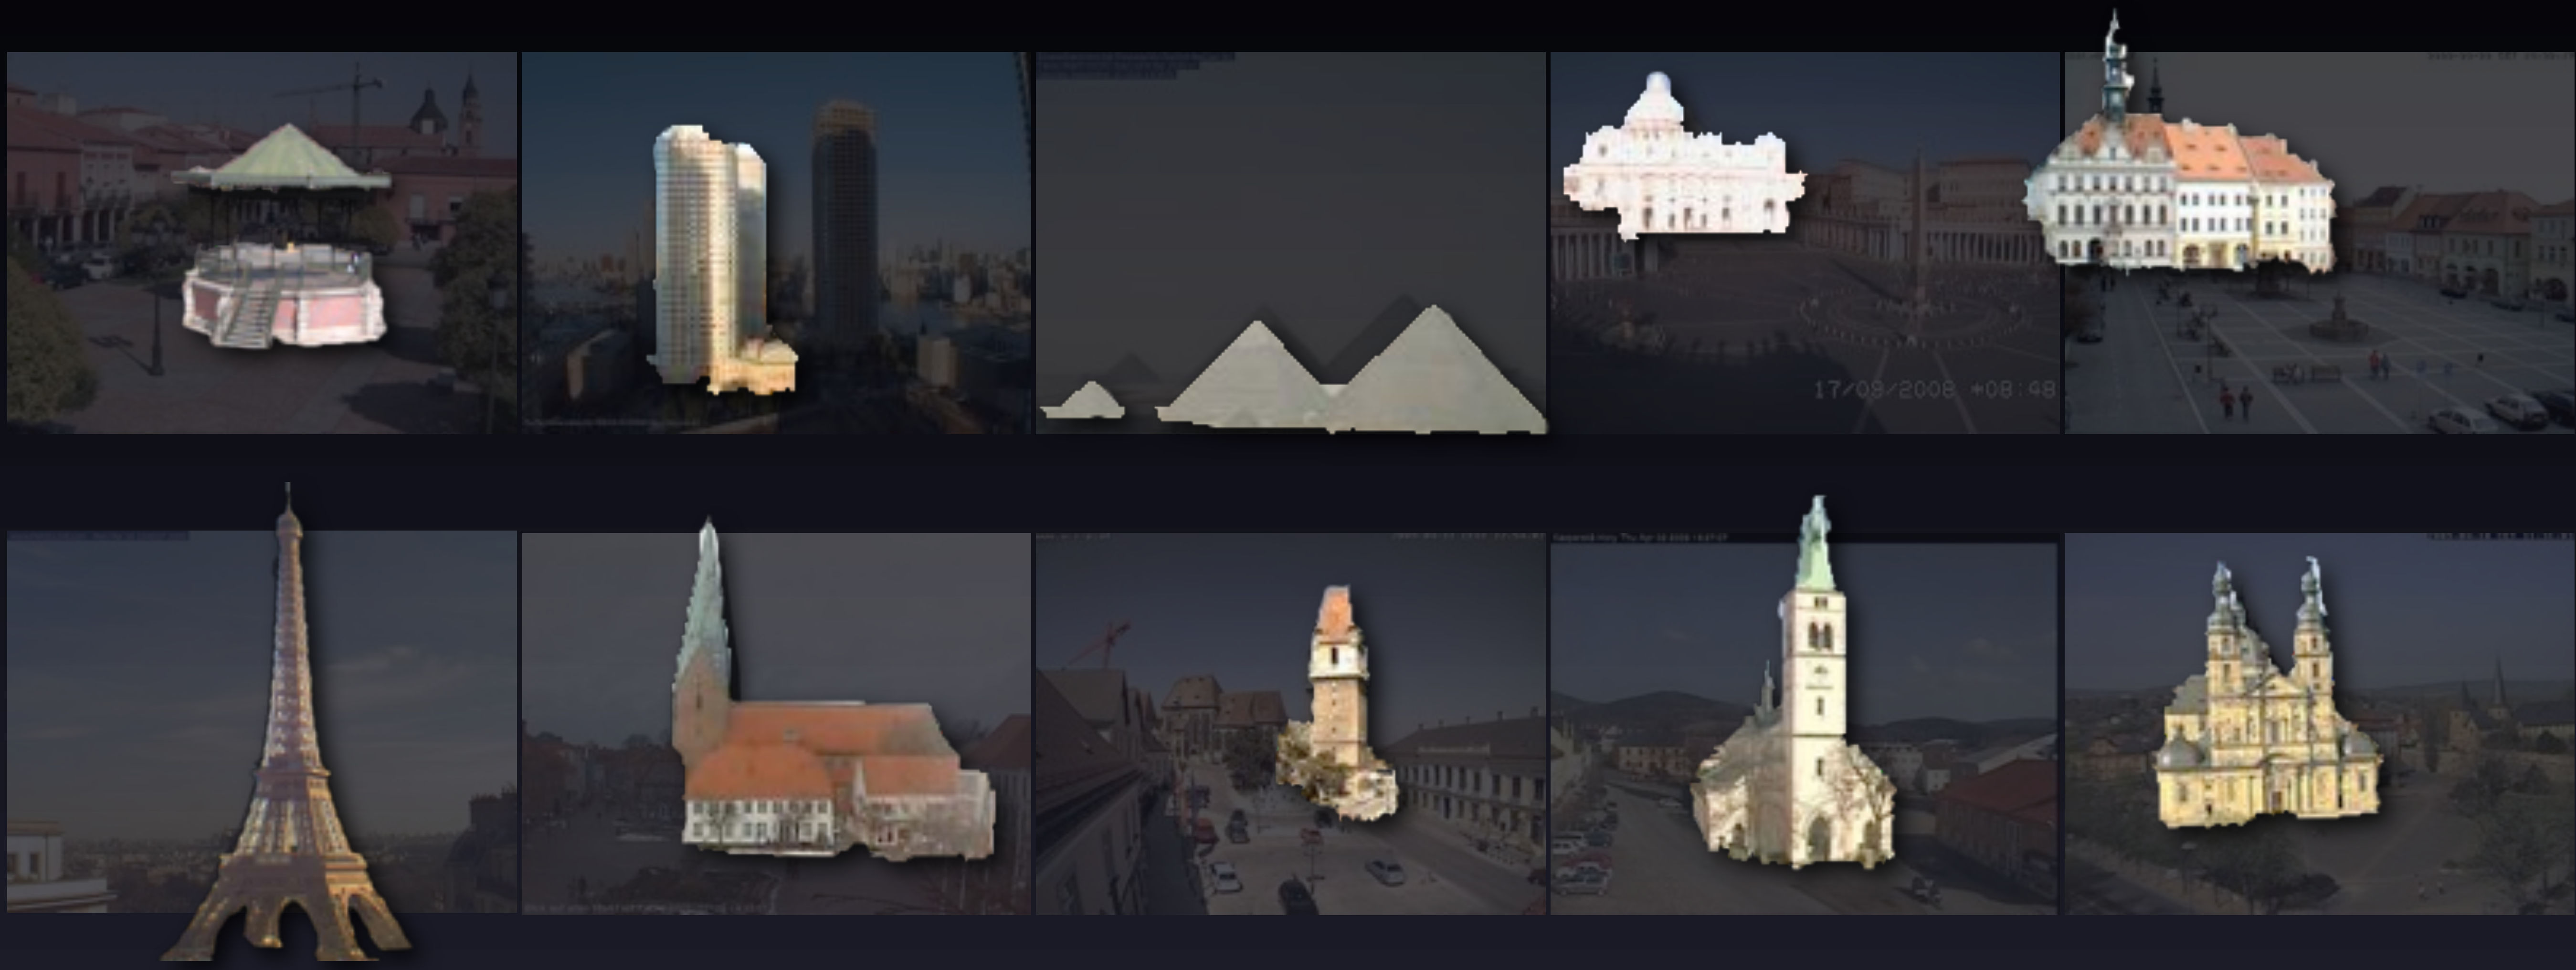

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Most webcams are not just empty scenes, but also contain interesting objects: architectural landmarks, buildings, etc... So it's also more than a database of scenes, but a database of objects, for which we have

Illumination-aware objects

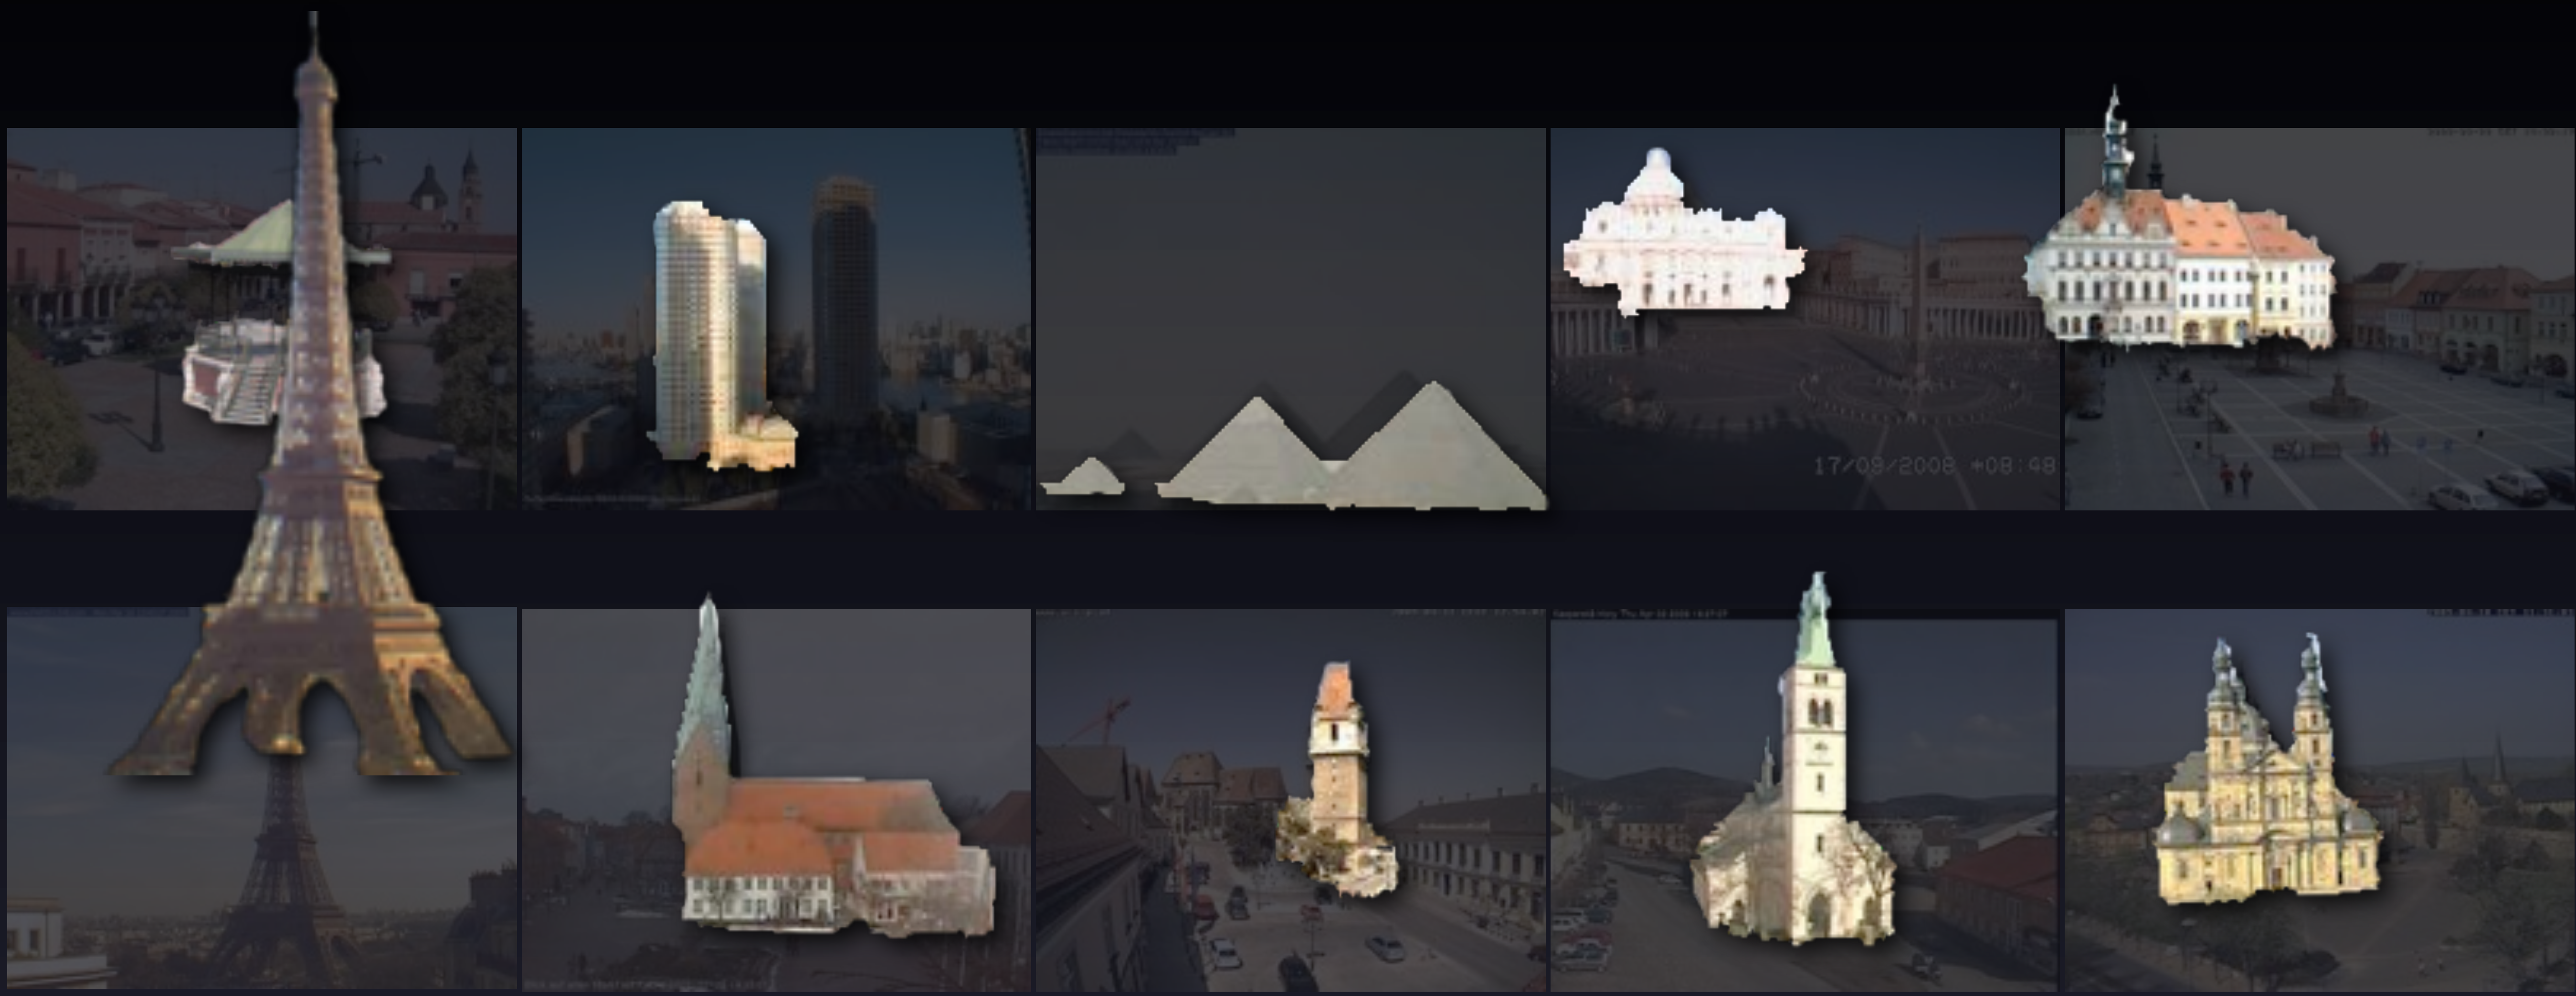

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Most webcams are not just empty scenes, but also contain interesting objects: architectural landmarks, buildings, etc... So it’s also more than a database of scenes, but a database of objects, for which we have

Illumination-aware objects

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

thousands of images of their appearance under varying lighting conditions, and the corresponding illumination parameters. So in this next section of my talk, I’ll show how we can use our webcam dataset to create a clipart library of objects that are aware of their illumination, and which can then be transferred to and from other webcams.

Illumination-aware objects

Object appearance

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

thousands of images of their appearance under varying lighting conditions, and the corresponding illumination parameters. So in this next section of my talk, I'll show how we can use our webcam dataset to create a clipart library of objects that are aware of their illumination, and which can then be transferred to and from other webcams.

Illumination-aware objects

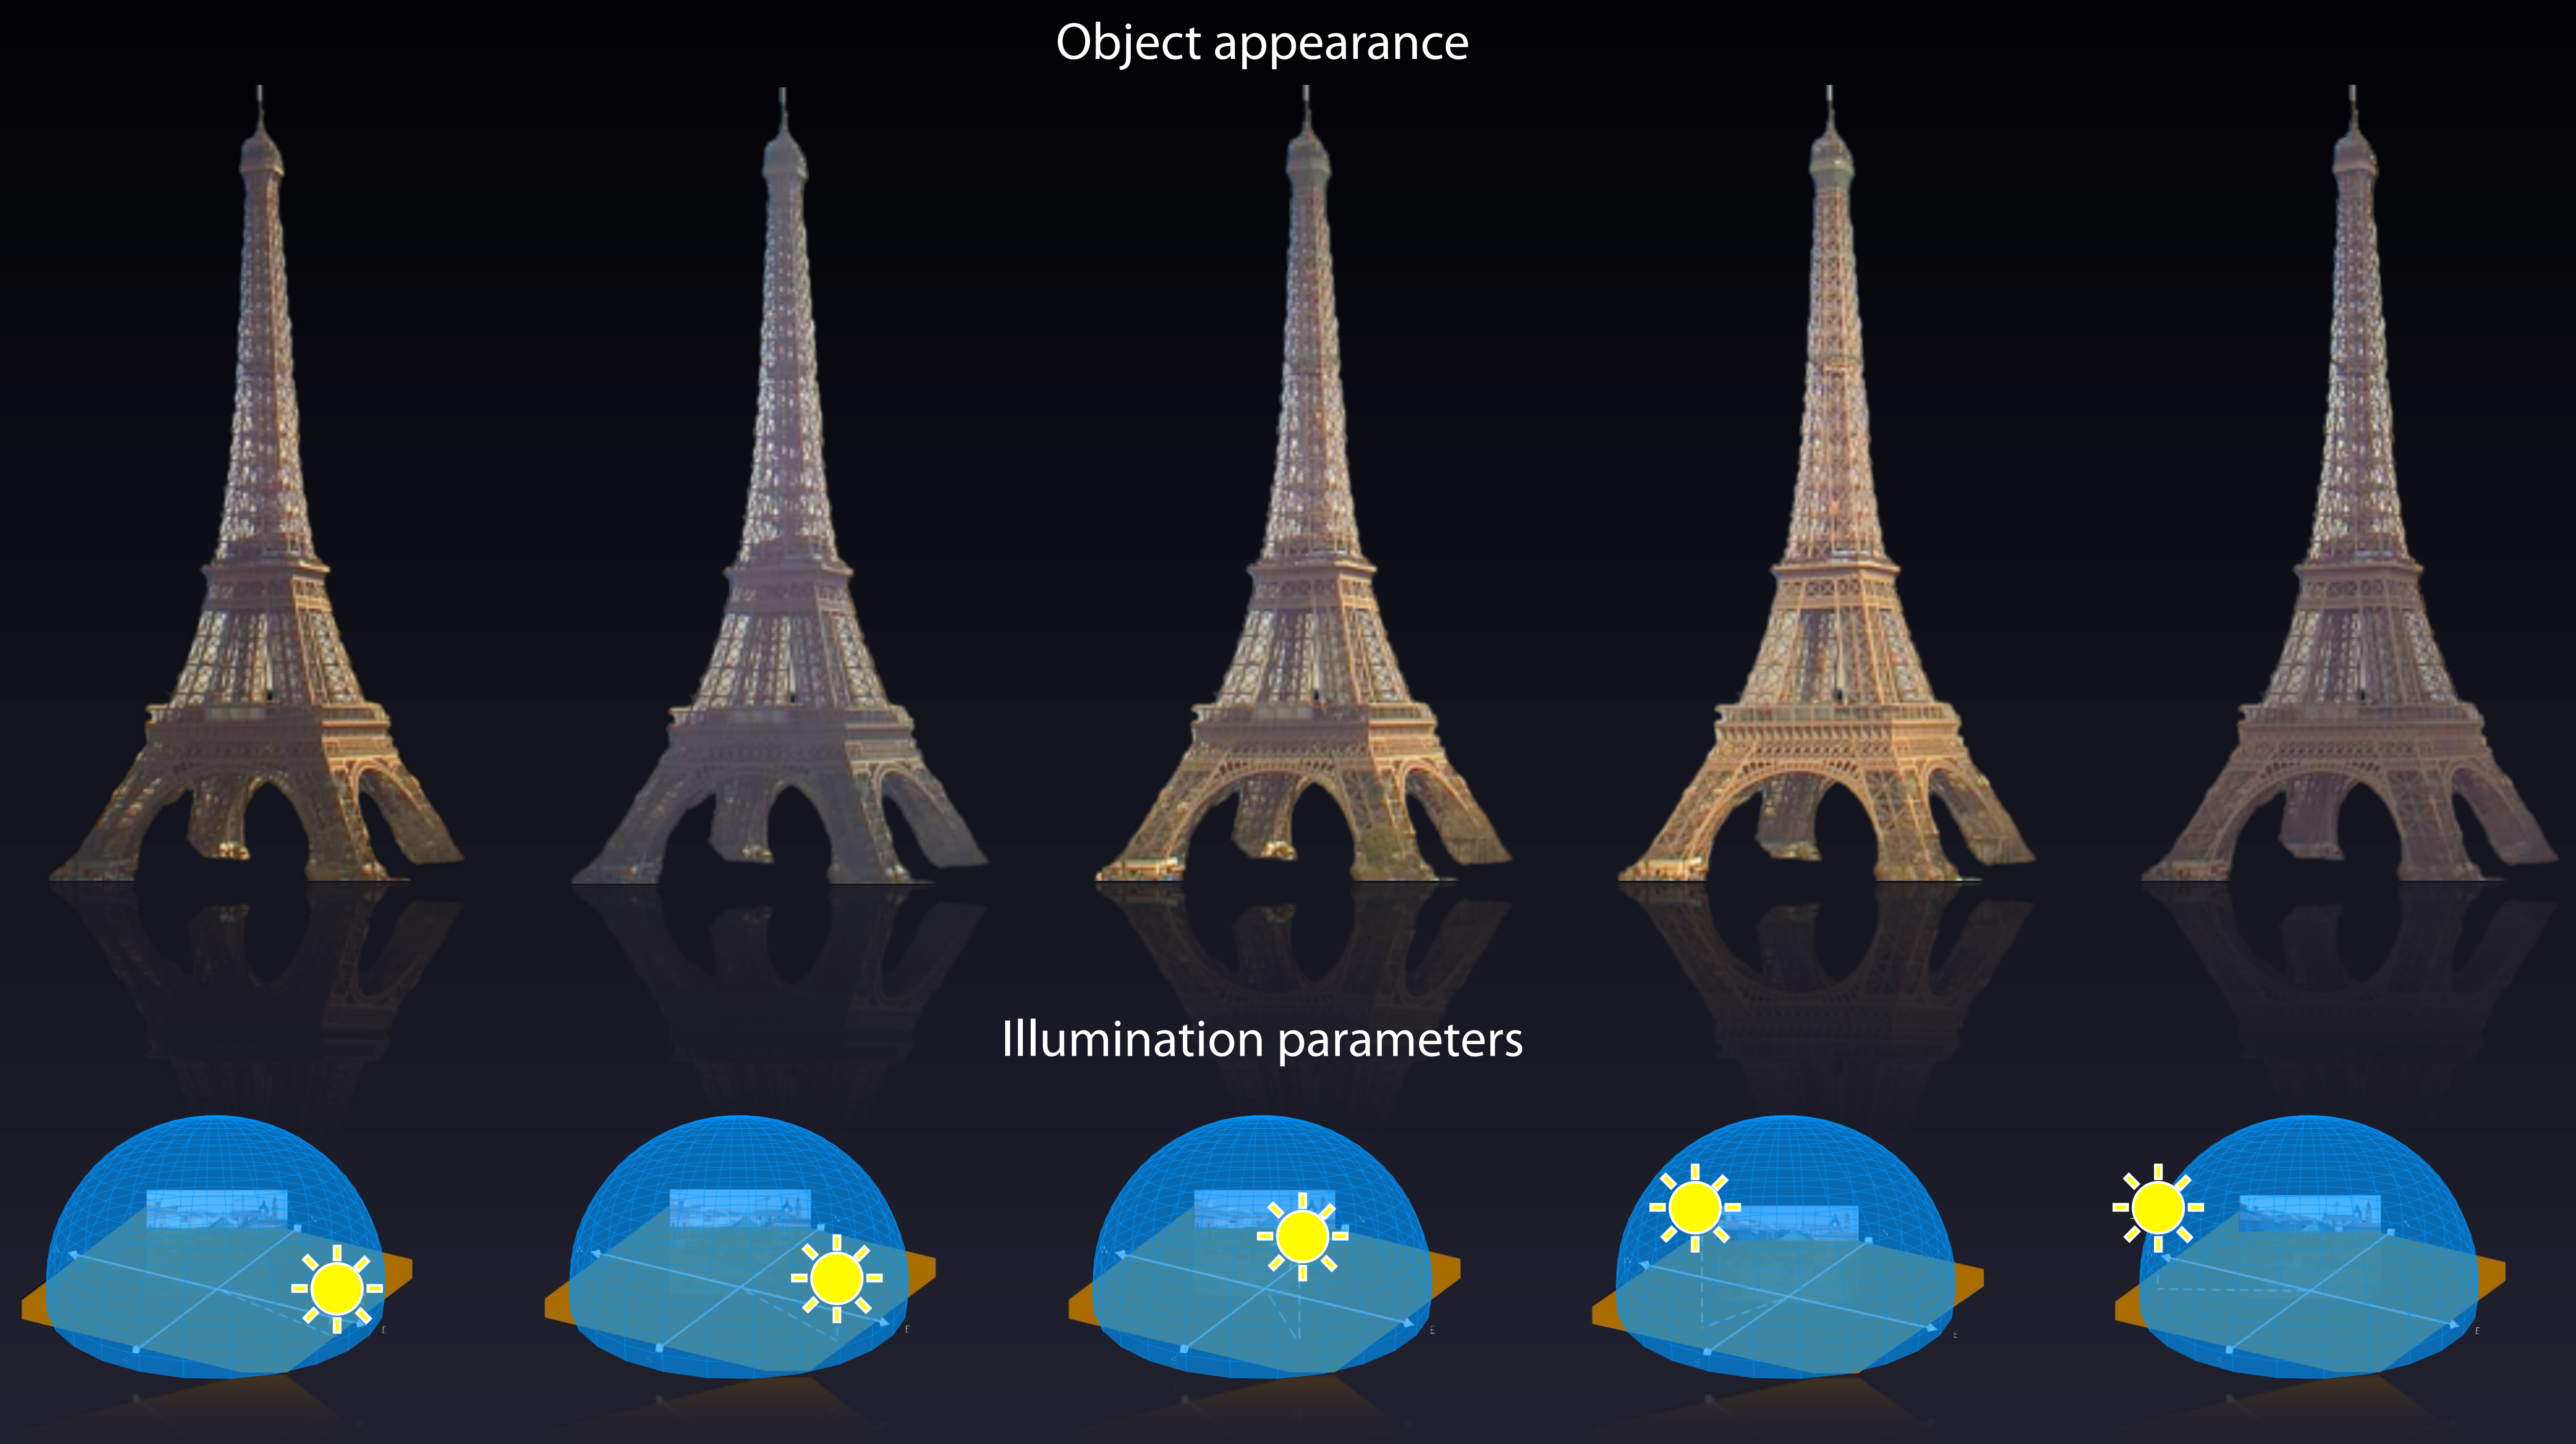

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

thousands of images of their appearance under varying lighting conditions, and the corresponding illumination parameters. So in this next section of my talk, I'll show how we can use our webcam dataset to create a clipart library of objects that are aware of their illumination, and which can then be transferred to and from other webcams.

Object transfer – Eiffel tower

Source

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

For instance, let’s say that we want to insert the Eiffel tower into the Berlin skyline. If we just copy and paste the pixels, the result looks wrong because of the widely different illumination conditions. Applying recoloring to match the global color distributions of the object with that of the background is very important, but it’s not enough to handle shading variations.

Object transfer – Eiffel tower

Source

Copy pixels

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

For instance, let's say that we want to insert the Eiffel tower into the Berlin skyline. If we just copy and paste the pixels, the result looks wrong because of the widely different illumination conditions. Applying recoloring to match the global color distributions of the object with that of the background is very important, but it's not enough to handle shading variations.

Object transfer – Eiffel tower

Source

Recoloring [Reinhard *et al.*, 2001]

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

For instance, let's say that we want to insert the Eiffel tower into the Berlin skyline. If we just copy and paste the pixels, the result looks wrong because of the widely different illumination conditions. Applying recoloring to match the global color distributions of the object with that of the background is very important, but it's not enough to handle shading variations.

Matching illumination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So instead, we can use our object database, and find an instance of the object with illumination conditions similar to our scene. The procedure is quite simple: compute the illumination parameters of each instance of the object, and of the scene, compute pairwise distances between each one, and keep the one with the smallest distance.

Matching illumination

Object

Scene

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So instead, we can use our object database, and find an instance of the object with illumination conditions similar to our scene. The procedure is quite simple: compute the illumination parameters of each instance of the object, and of the scene, compute pairwise distances between each one, and keep the one with the smallest distance.

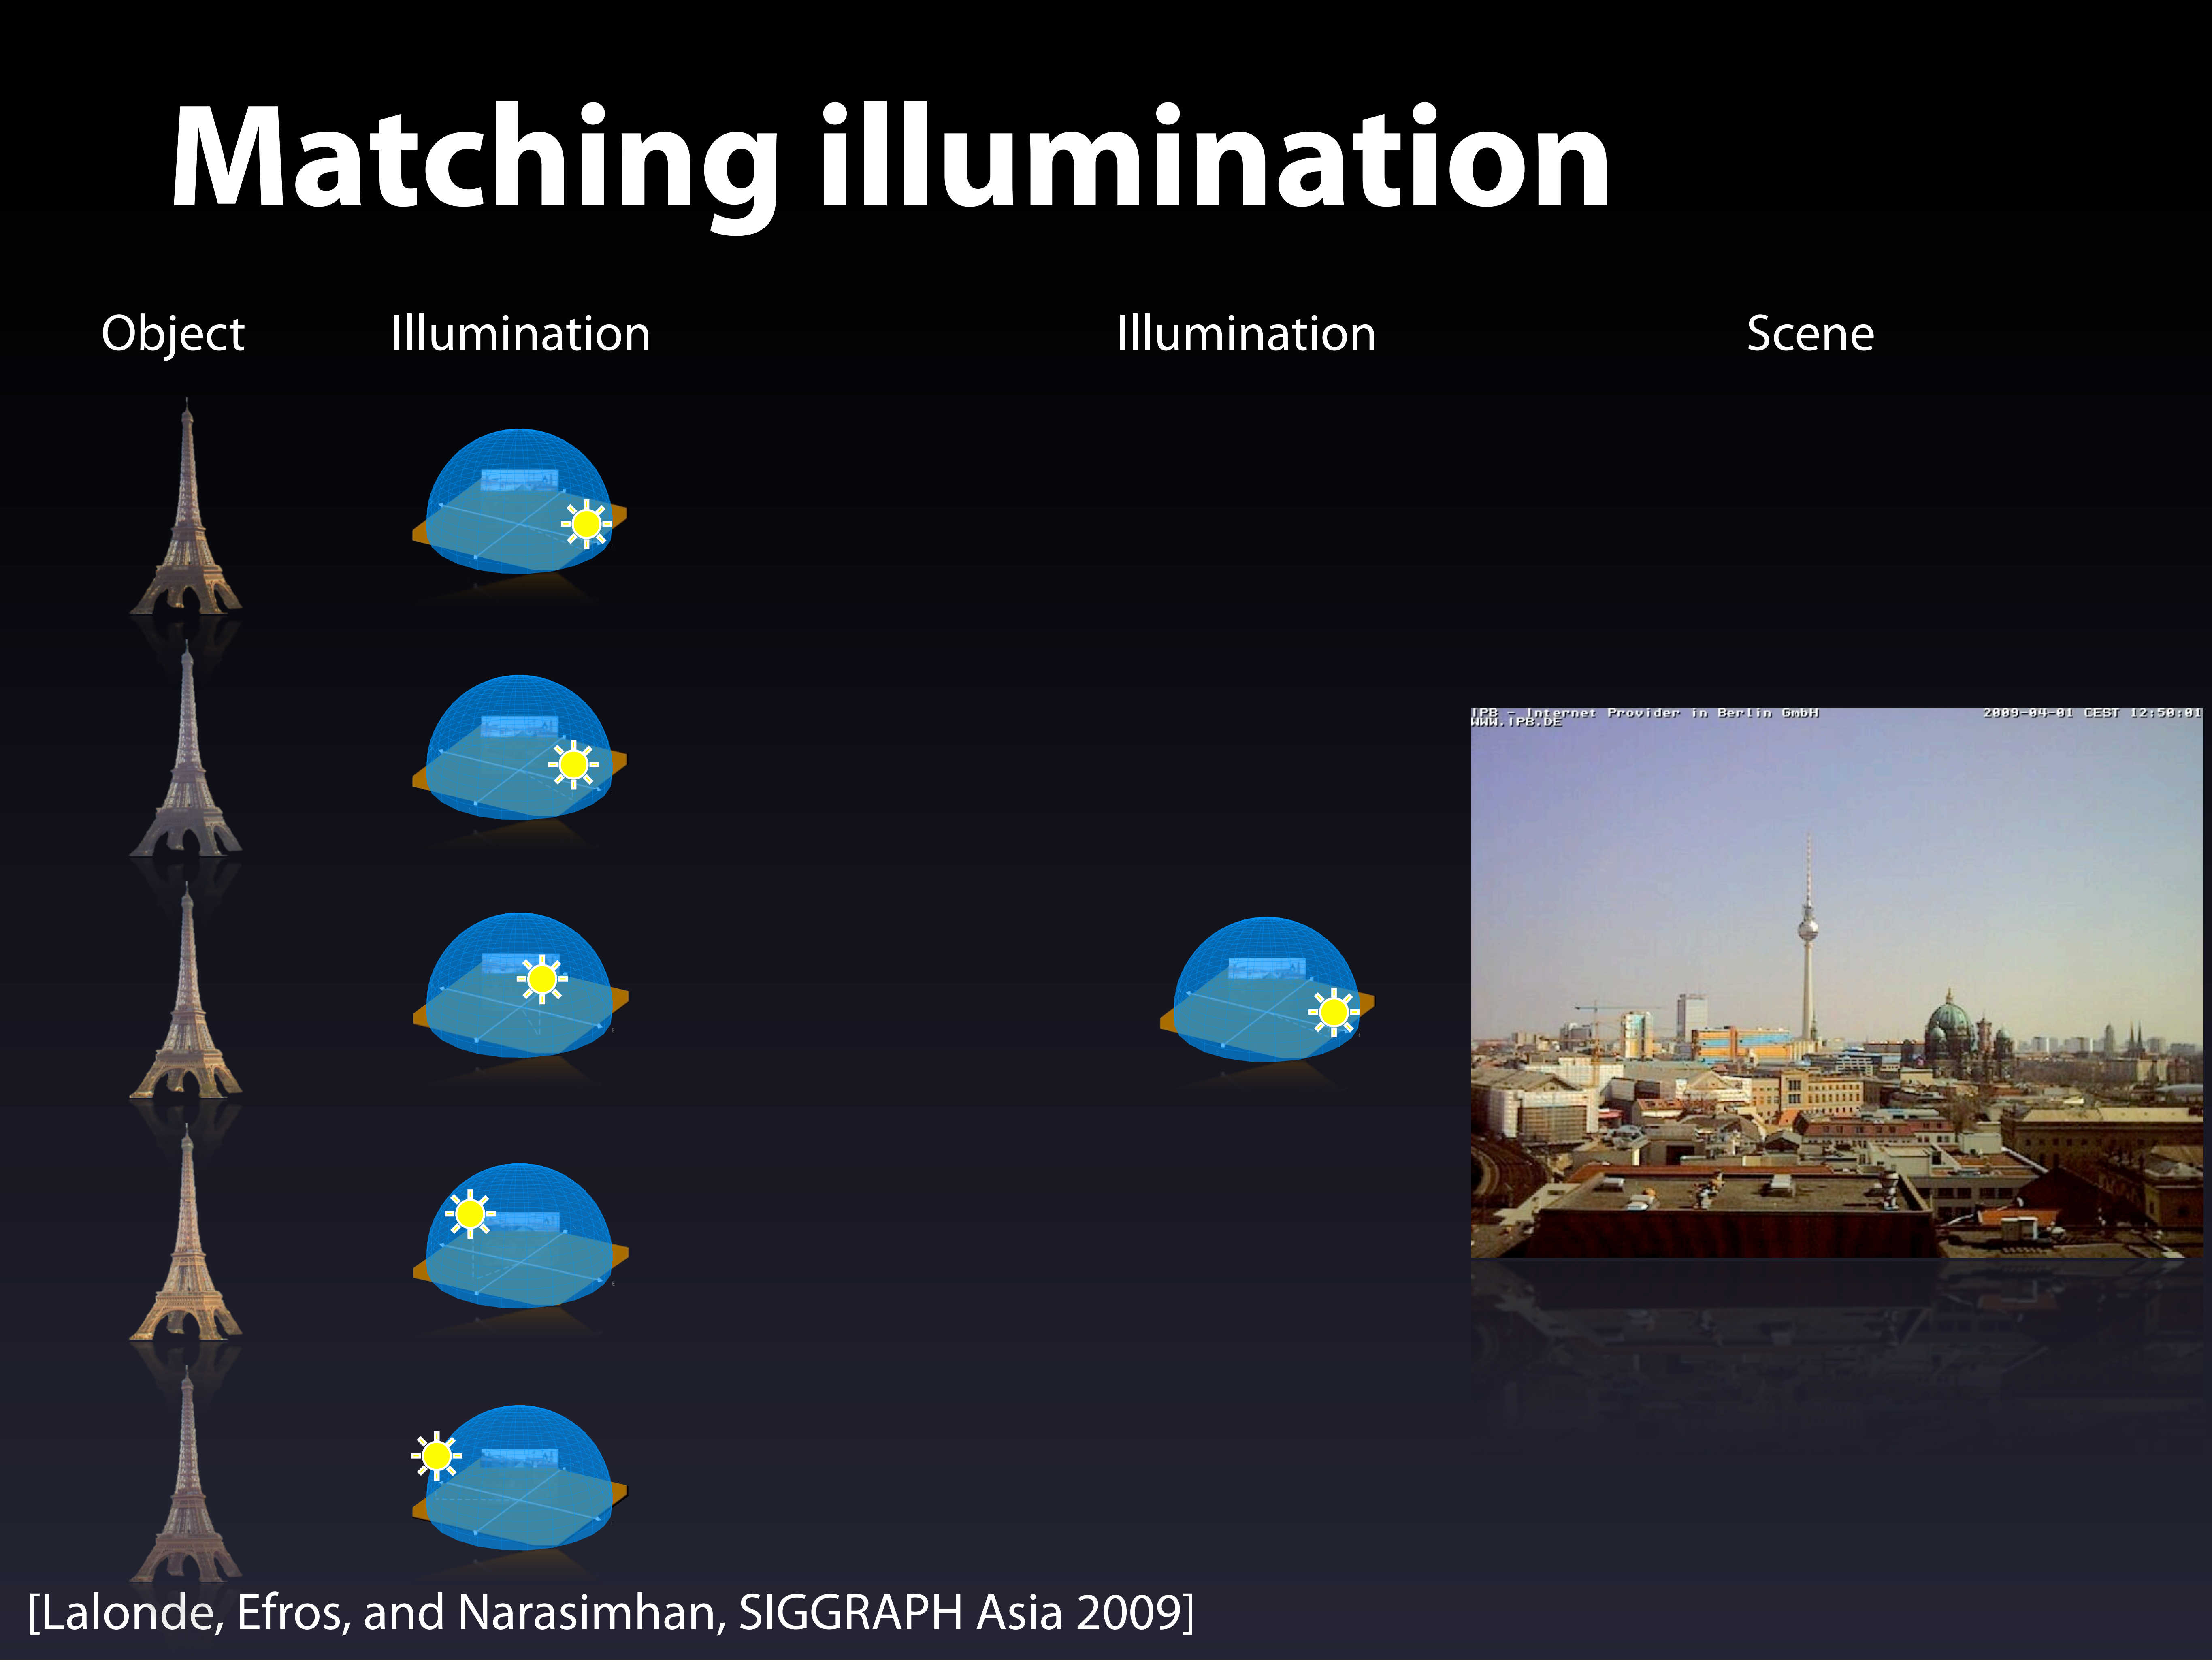

So instead, we can use our object database, and find an instance of the object with illumination conditions similar to our scene. The procedure is quite simple: compute the illumination parameters of each instance of the object, and of the scene, compute pairwise distances between each one, and keep the one with the smallest distance.

Matching illumination

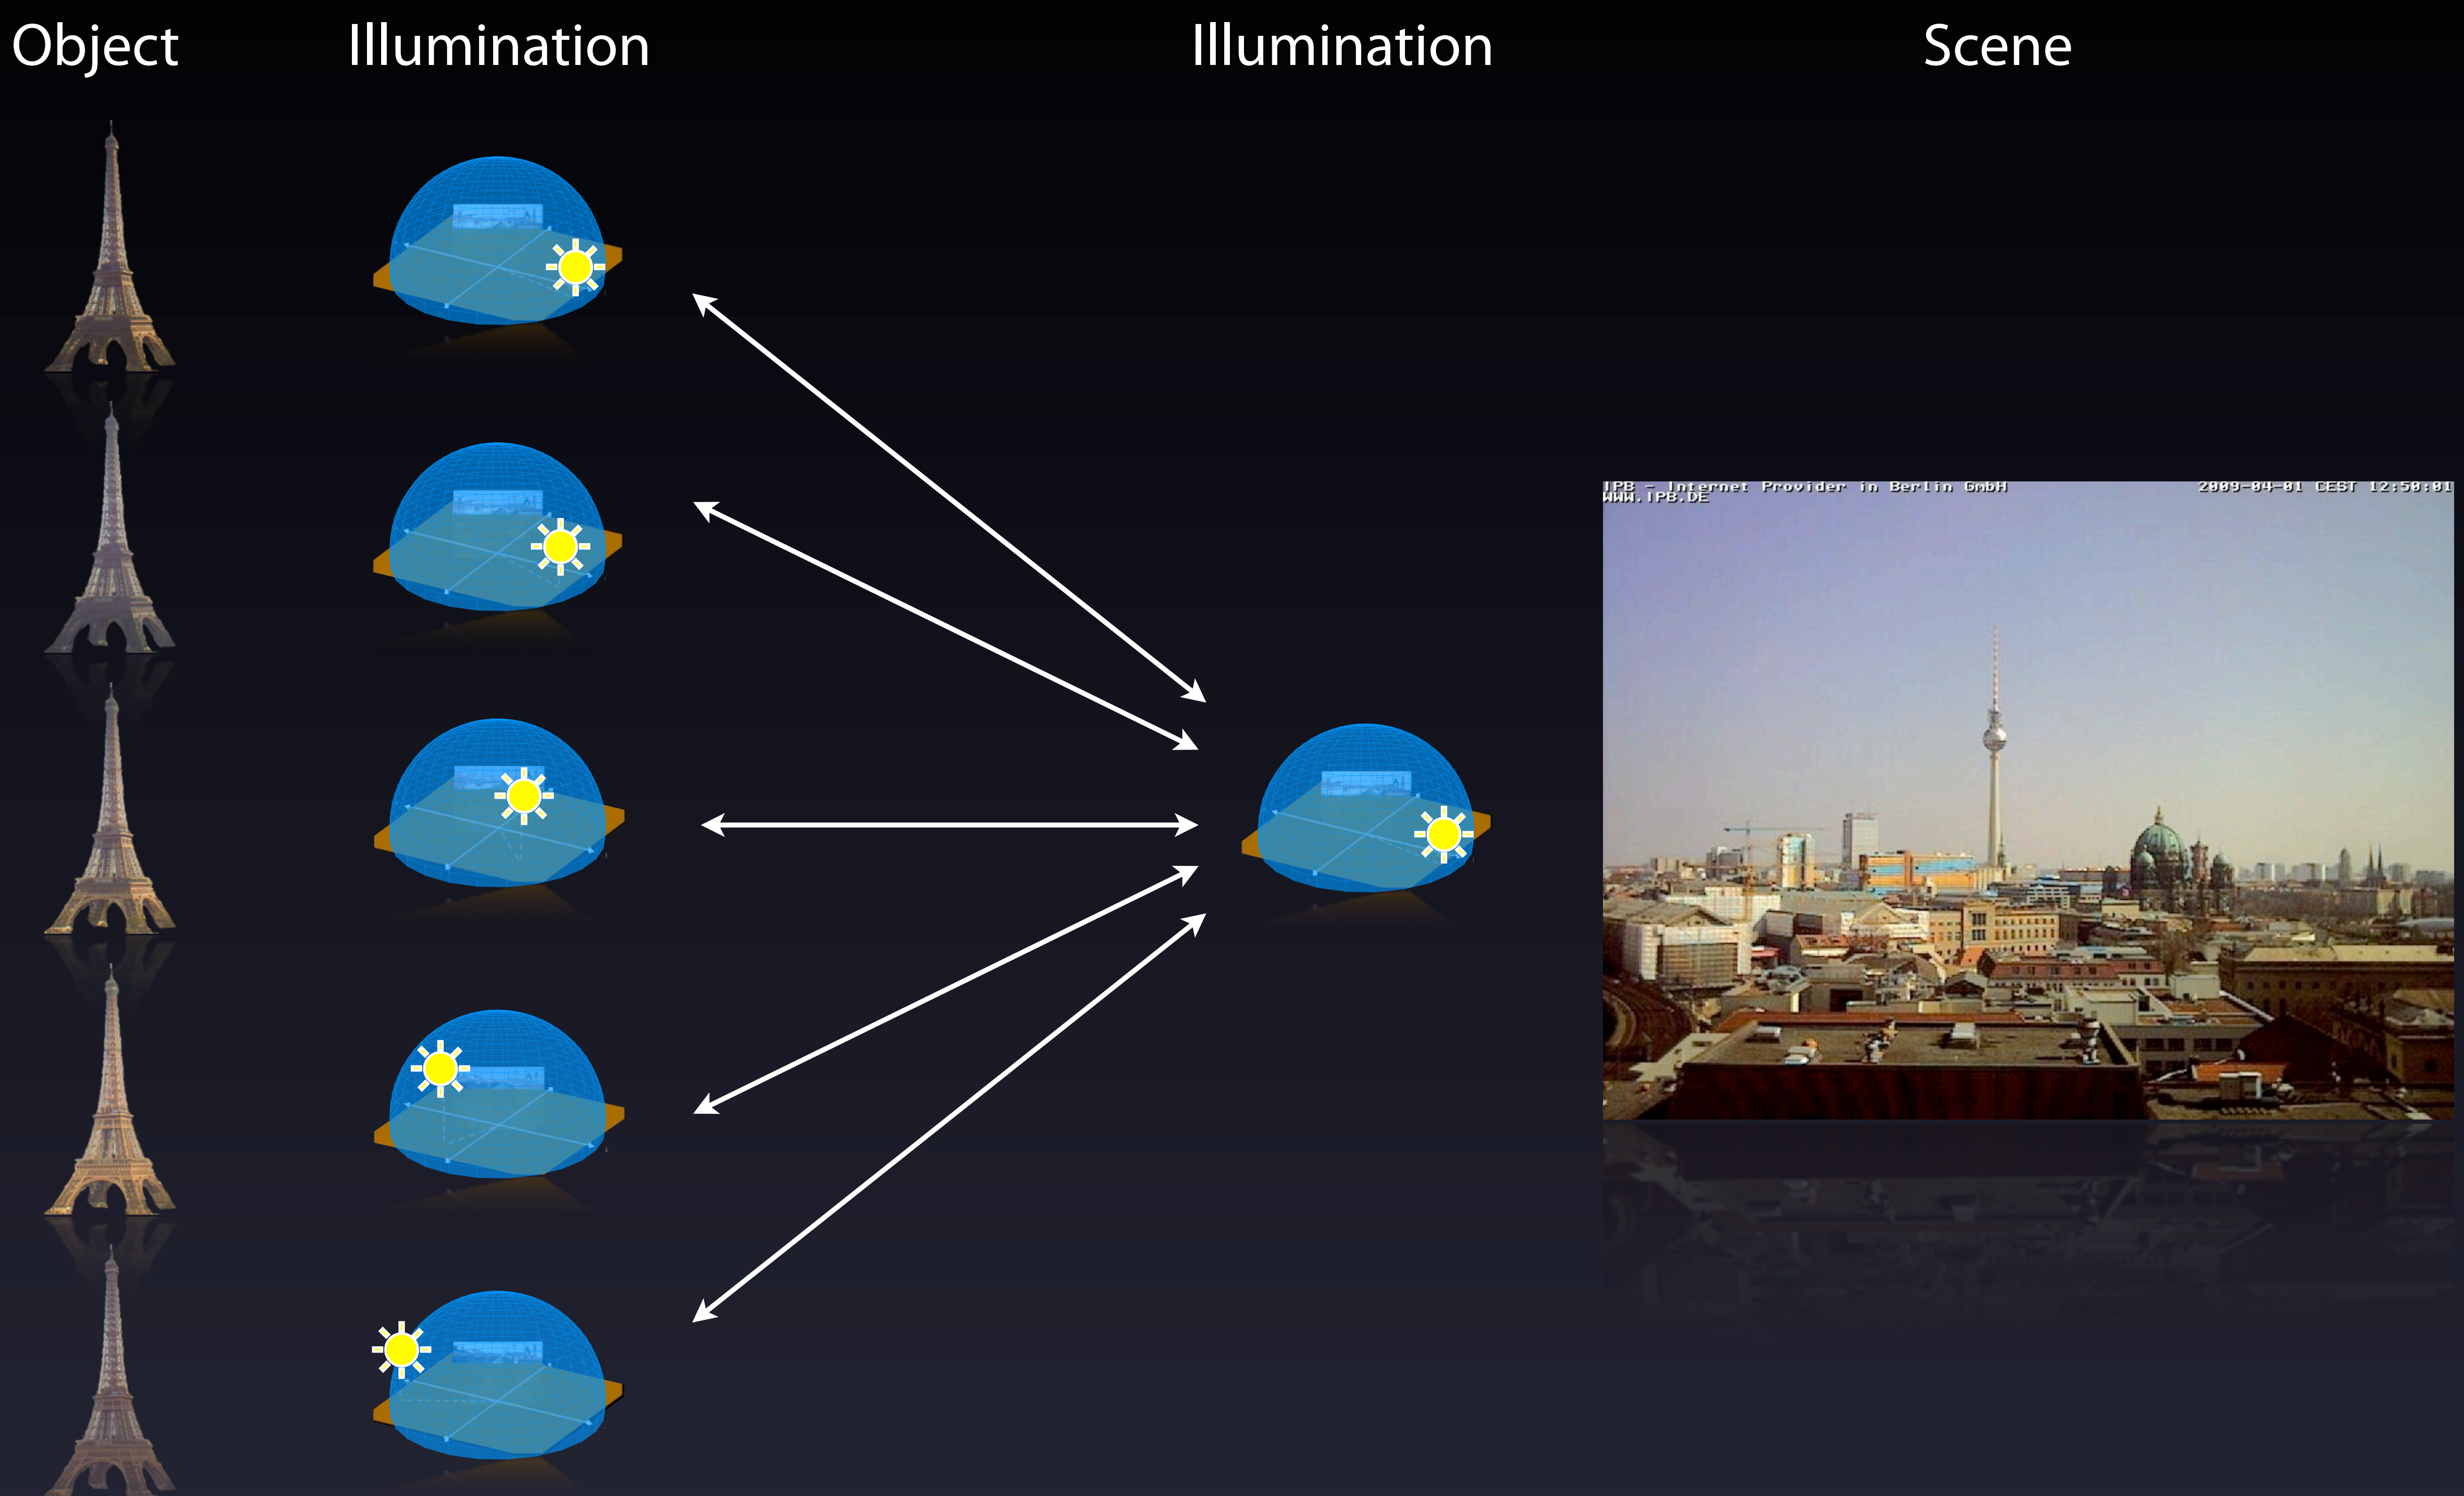

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So instead, we can use our object database, and find an instance of the object with illumination conditions similar to our scene. The procedure is quite simple: compute the illumination parameters of each instance of the object, and of the scene, compute pairwise distances between each one, and keep the one with the smallest distance.

Matching illumination

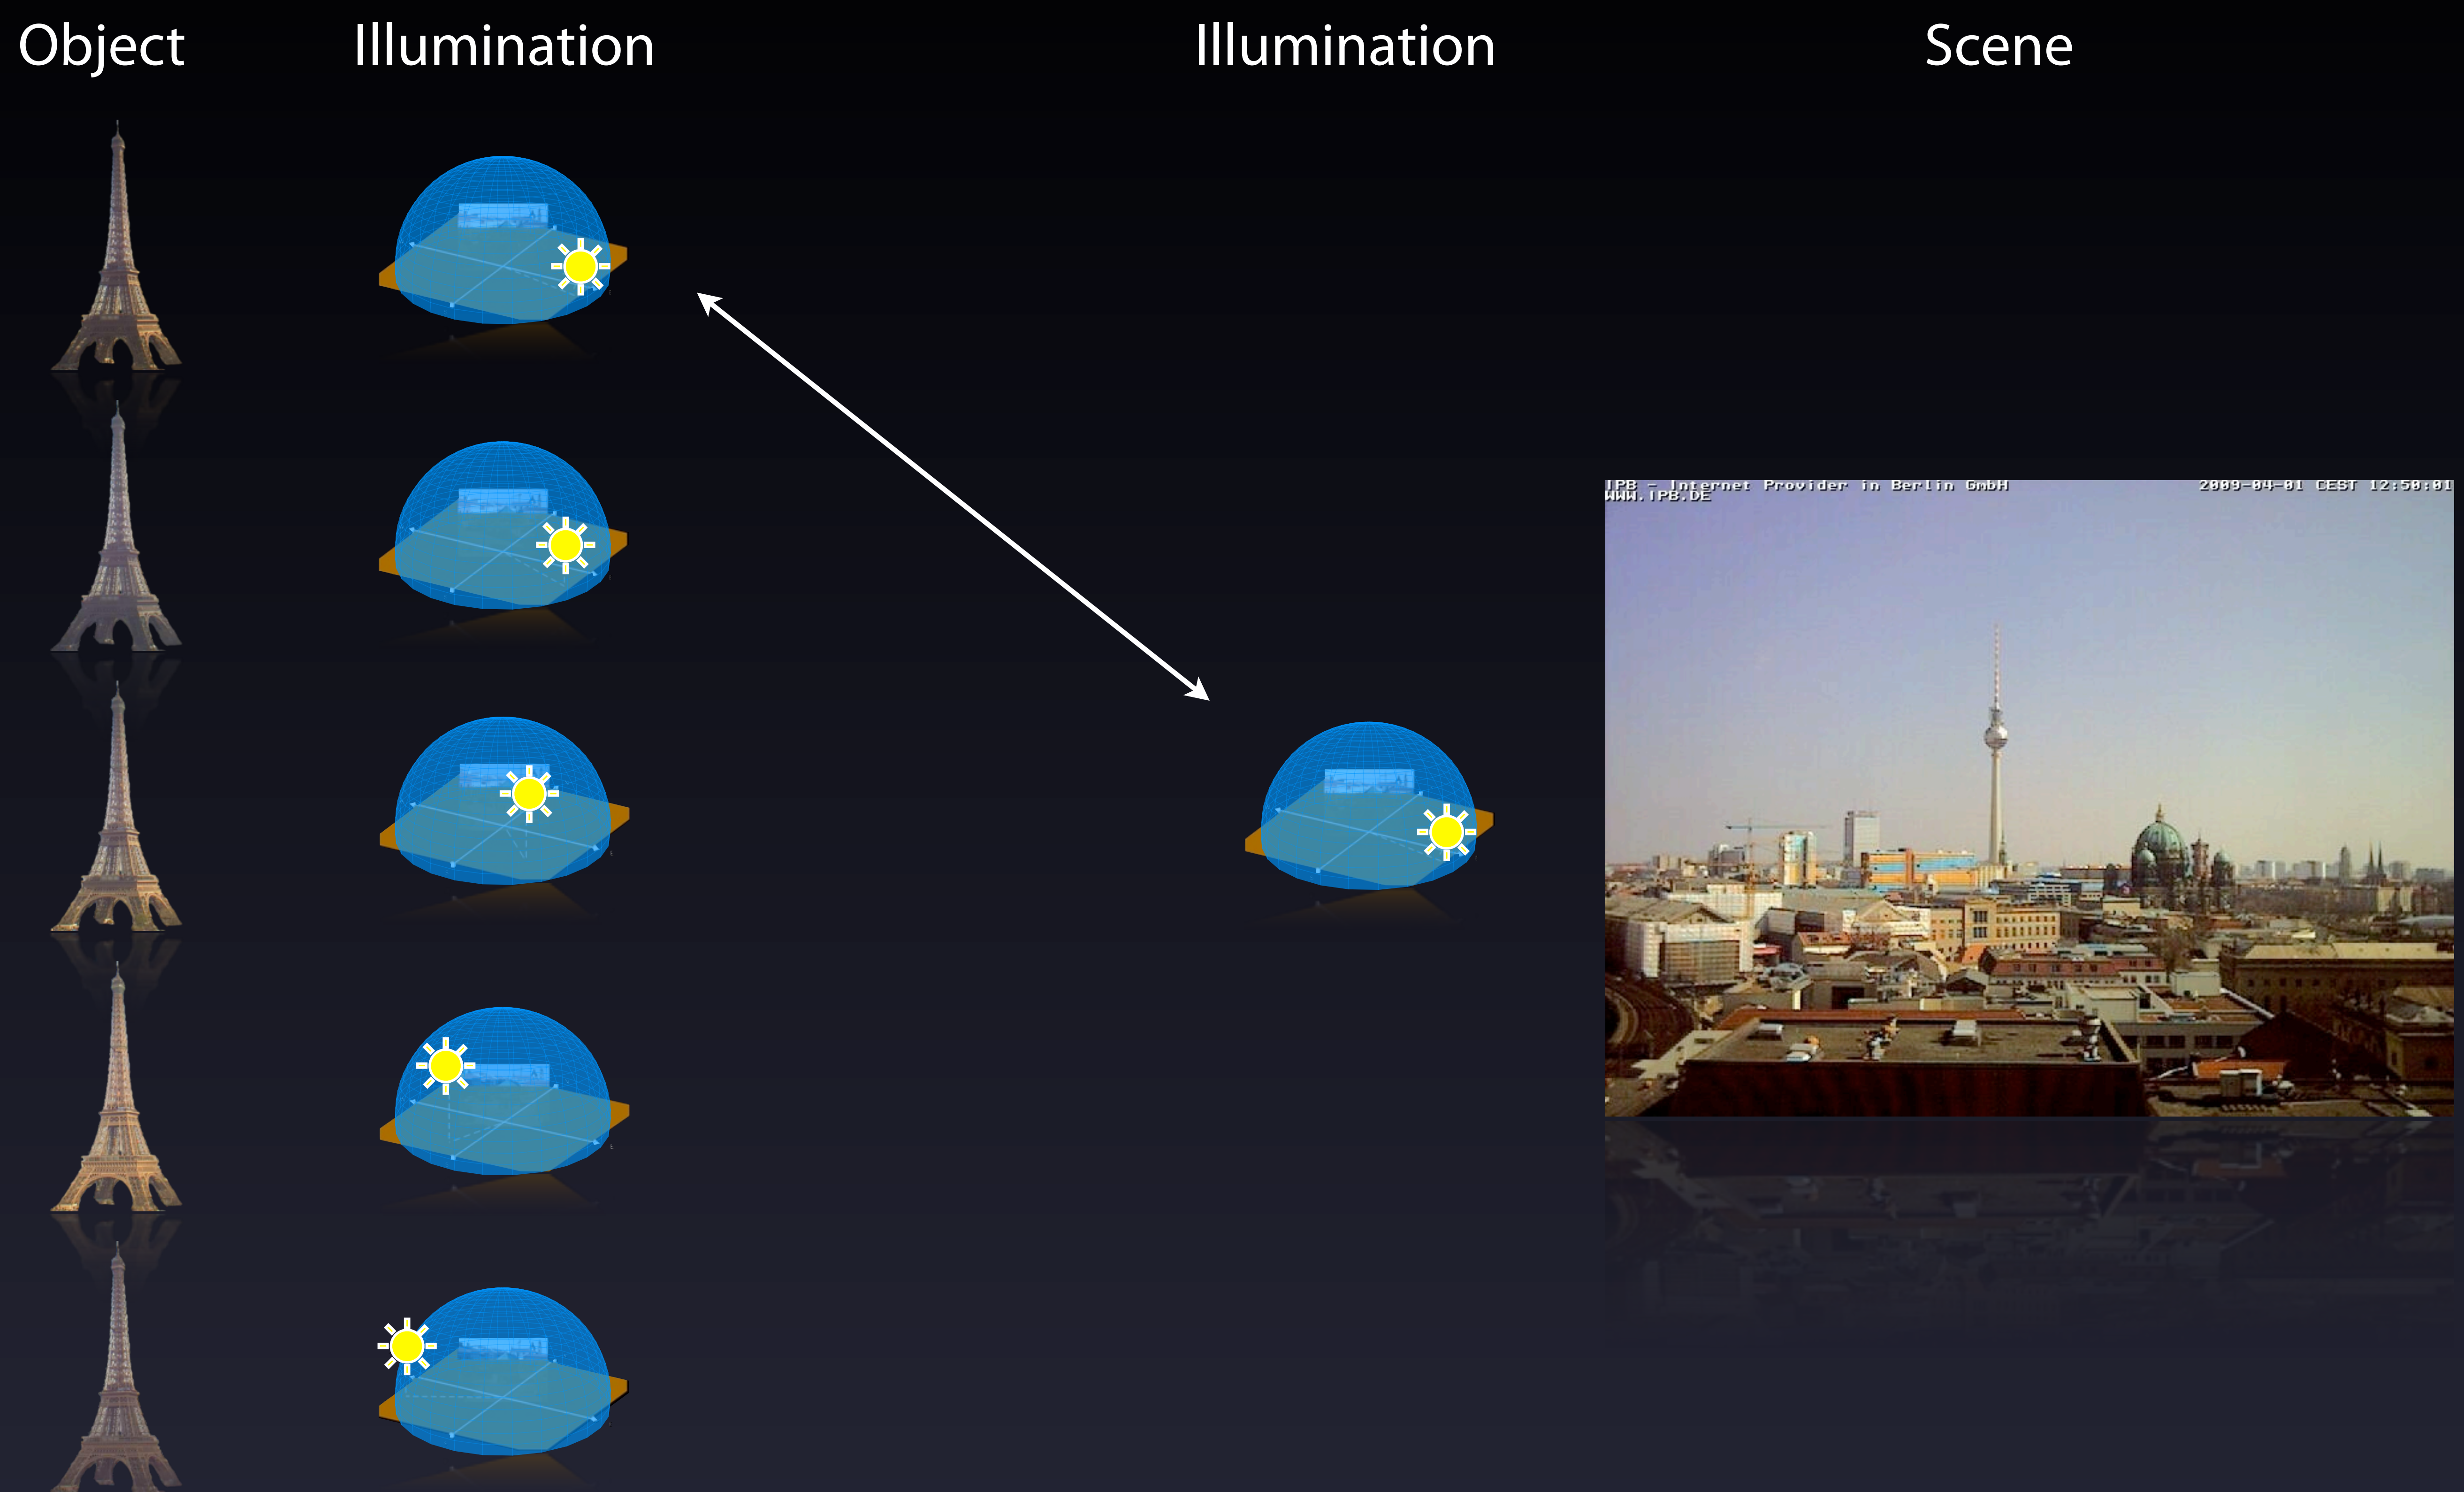

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

So instead, we can use our object database, and find an instance of the object with illumination conditions similar to our scene. The procedure is quite simple: compute the illumination parameters of each instance of the object, and of the scene, compute pairwise distances between each one, and keep the one with the smallest distance.

Object transfer – Eiffel tower

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Doing so, we get a result which is consistent with the illumination of the destination image, in that example, the sun is shining on the scene, and it comes from the right.

Object transfer – Eiffel tower

Object with similar illumination

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Doing so, we get a result which is consistent with the illumination of the destination image, in that example, the sun is shining on the scene, and it comes from the right.

Object transfer – Eiffel tower

Object with similar illumination

Our result

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Doing so, we get a result which is consistent with the illumination of the destination image, in that example, the sun is shining on the scene, and it comes from the right.

Object transfer – Eiffel tower

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

And we can do this for different times of day. Notice how the shading on the tower matches the rest of the scene.

Object transfer – Eiffel tower

19:00

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

And we can do this for different times of day. Notice how the shading on the tower matches the rest of the scene.

Object transfer – buildings

Source

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

In this other example, we wish to transfer these buildings from Tokyo into Berkeley. Again, simply copying the pixels is not enough, but we need to find an instance of the object under similar lighting conditions.

Object transfer – buildings

Source

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

In this other example, we wish to transfer these buildings from Tokyo into Berkeley. Again, simply copying the pixels is not enough, but we need to find an instance of the object under similar lighting conditions.

Object transfer – buildings

Source

Transfer result

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

In this other example, we wish to transfer these buildings from Tokyo into Berkeley. Again, simply copying the pixels is not enough, but we need to find an instance of the object under similar lighting conditions.

Object transfer – buildings

Transfer result: reflection

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Here's another result, showing that we can accurately replicate complex reflection effects, like the sun in the building windows.

Object transfer – buildings

Transfer result: reflection

Transfer result: sunset

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Here's another result, showing that we can accurately replicate complex reflection effects, like the sun in the building windows.

Limitations

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

The main limitation of our work at this point is in the interaction between the object and its new background, especially shadows. In the 2-D object transfer case, we currently transfer the object alone without its shadow. Since the objects are usually fairly large, they often cast shadows on the surrounding scene, which might be arbitrarily complex, making it very hard to detect and transfer accurately.

Limitations

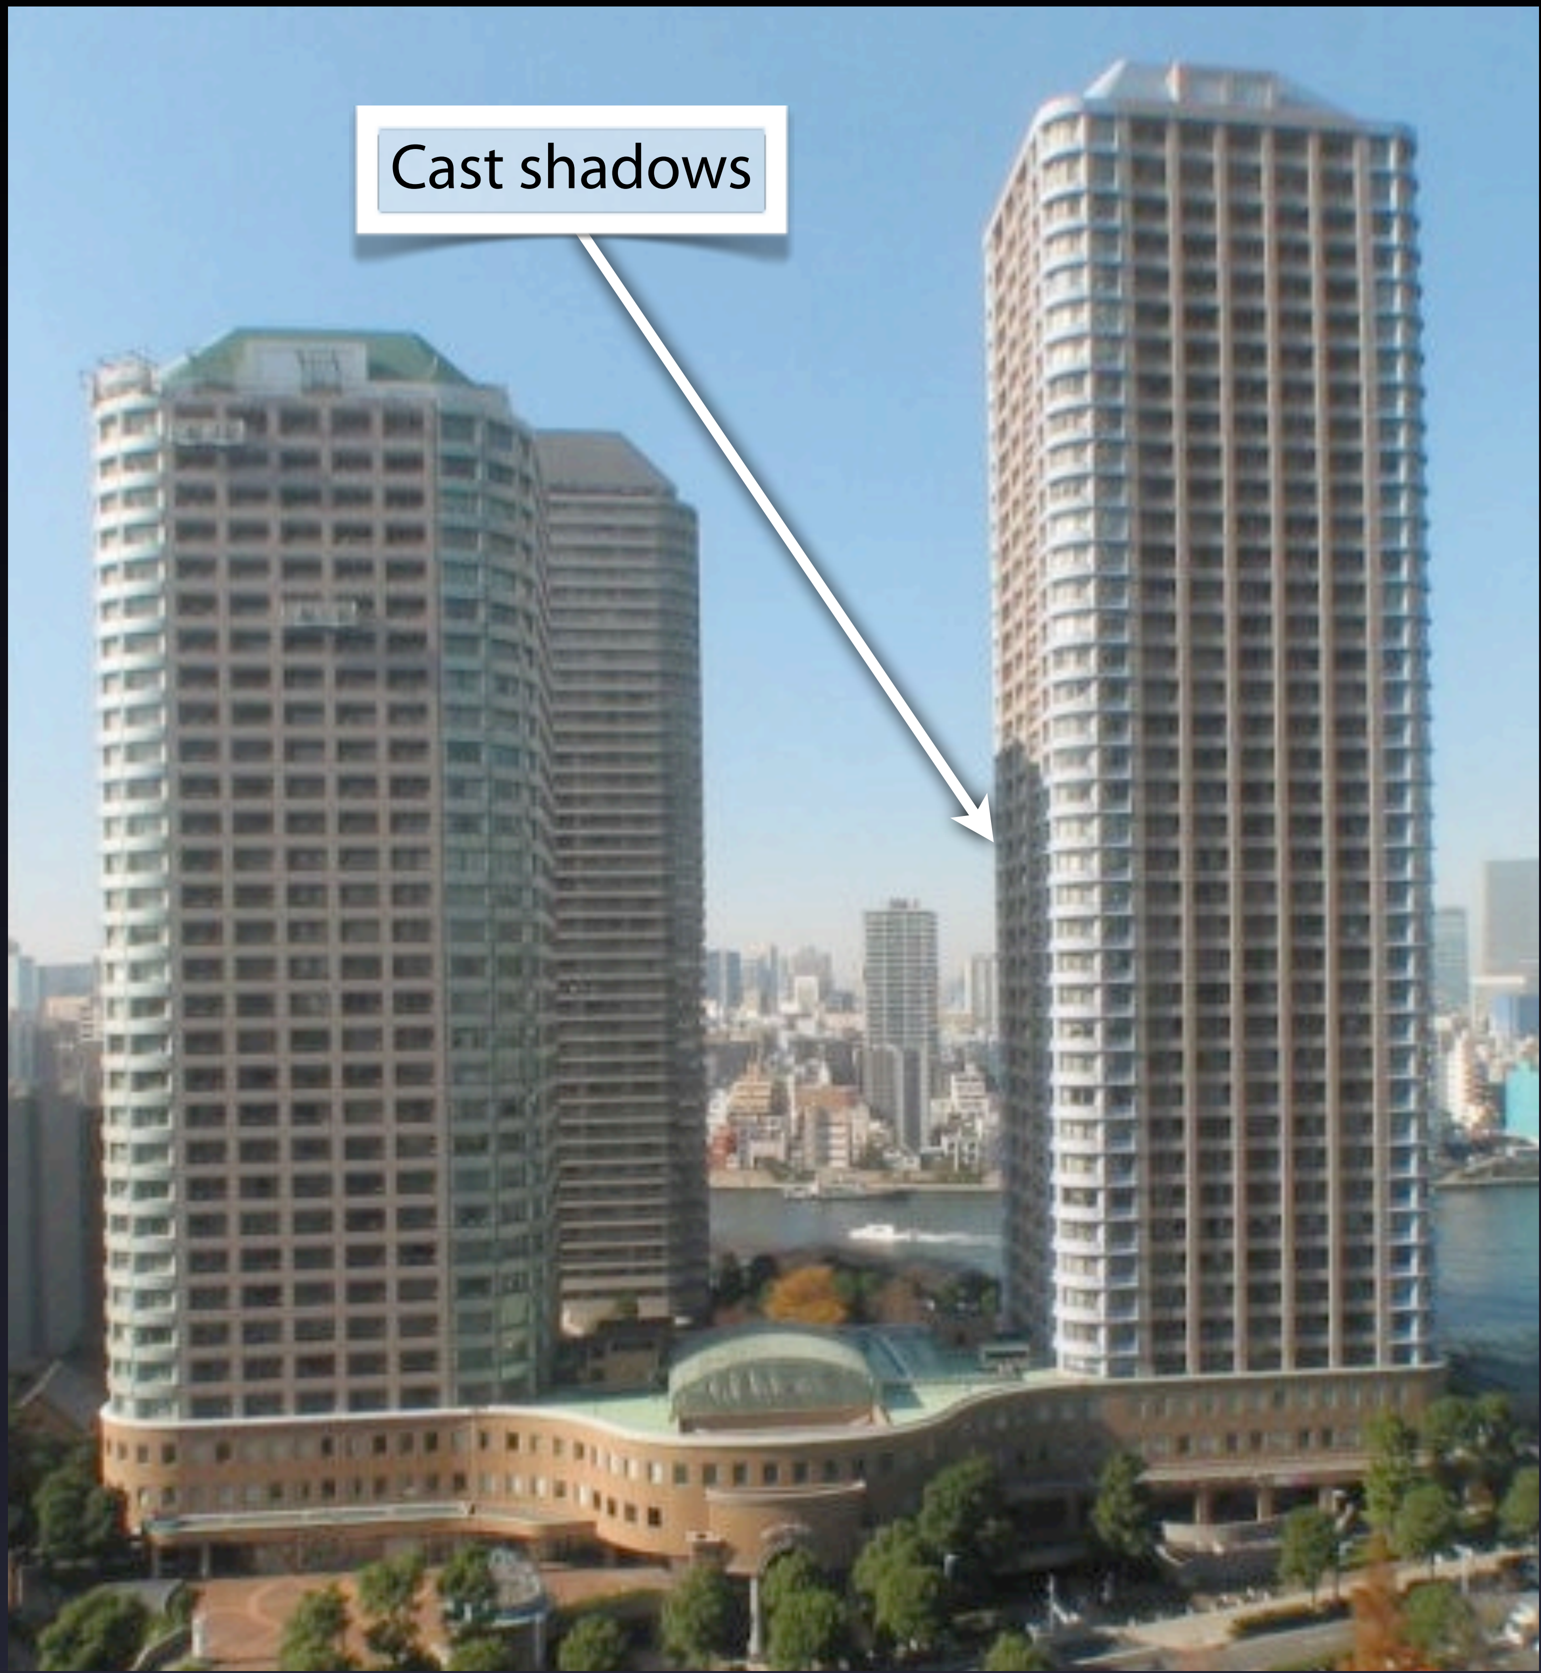

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

The main limitation of our work at this point is in the interaction between the object and its new background, especially shadows. In the 2-D object transfer case, we currently transfer the object alone without its shadow. Since the objects are usually fairly large, they often cast shadows on the surrounding scene, which might be arbitrarily complex, making it very hard to detect and transfer accurately.

Limitations

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Shadows are easier to handle when inserting 3-D objects, because the rendering engine can generate them accurately. Right now our scene model is a simple ground plane surrounding the object, so that does not cast shadows on other objects, but all that requires is a more complex model of the geometry. Perhaps one way to handle these limitations is to use existing techniques which try to estimate geometry from webcams.

Limitations

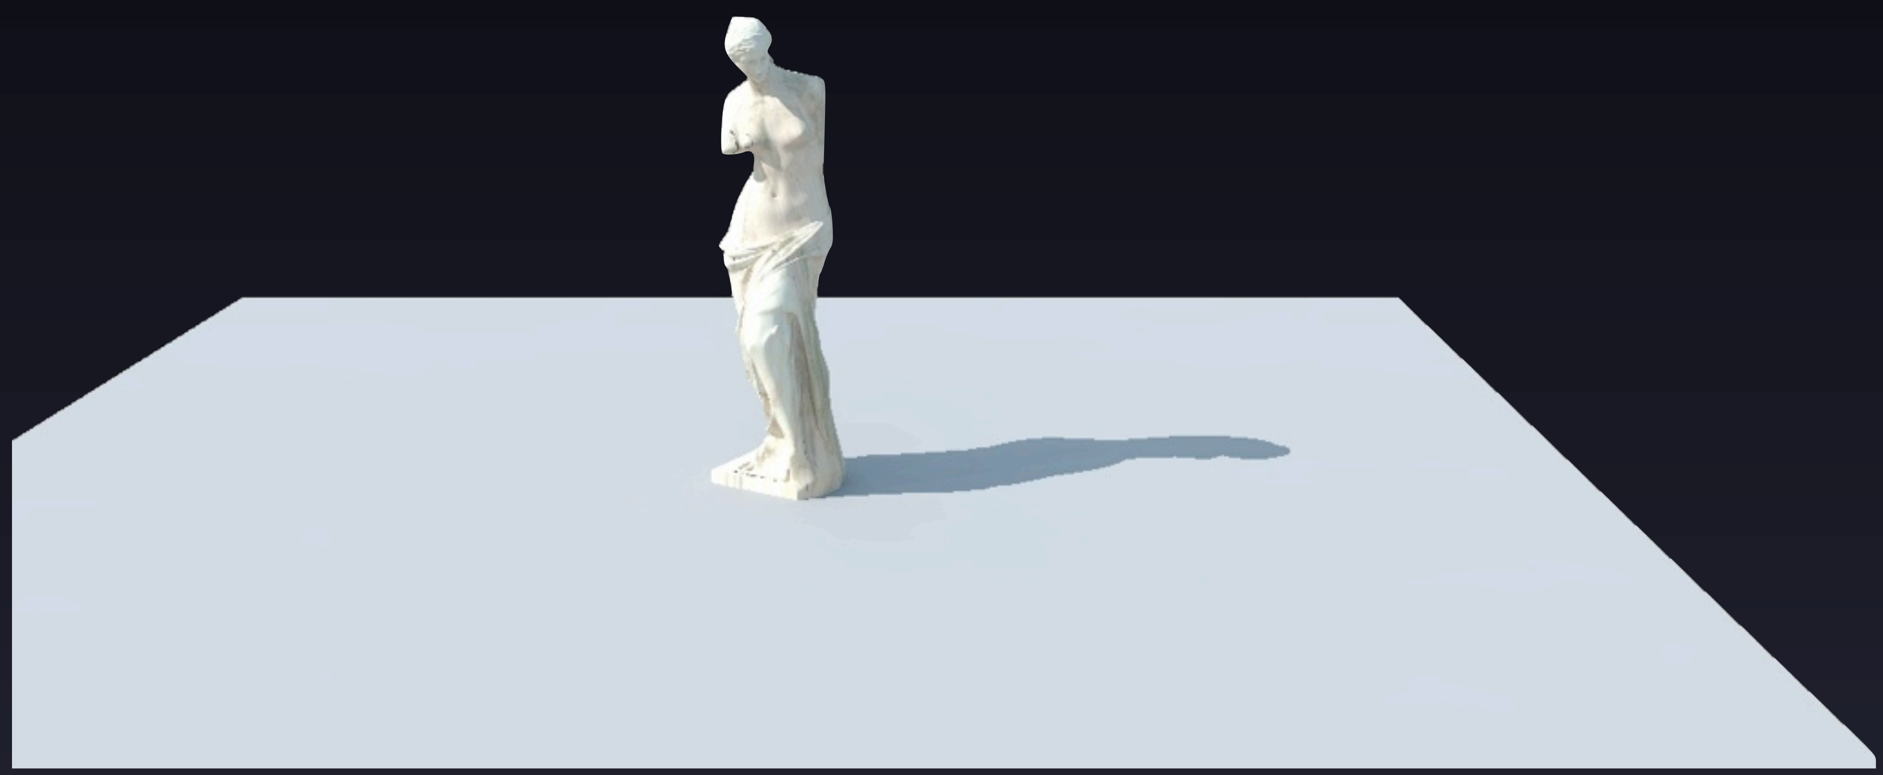

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Shadows are easier to handle when inserting 3-D objects, because the rendering engine can generate them accurately. Right now our scene model is a simple ground plane surrounding the object, so that does not cast shadows on other objects, but all that requires is a more complex model of the geometry. Perhaps one way to handle these limitations is to use existing techniques which try to estimate geometry from webcams.

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

To conclude, we’ve presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it’s a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

Database of image sequences

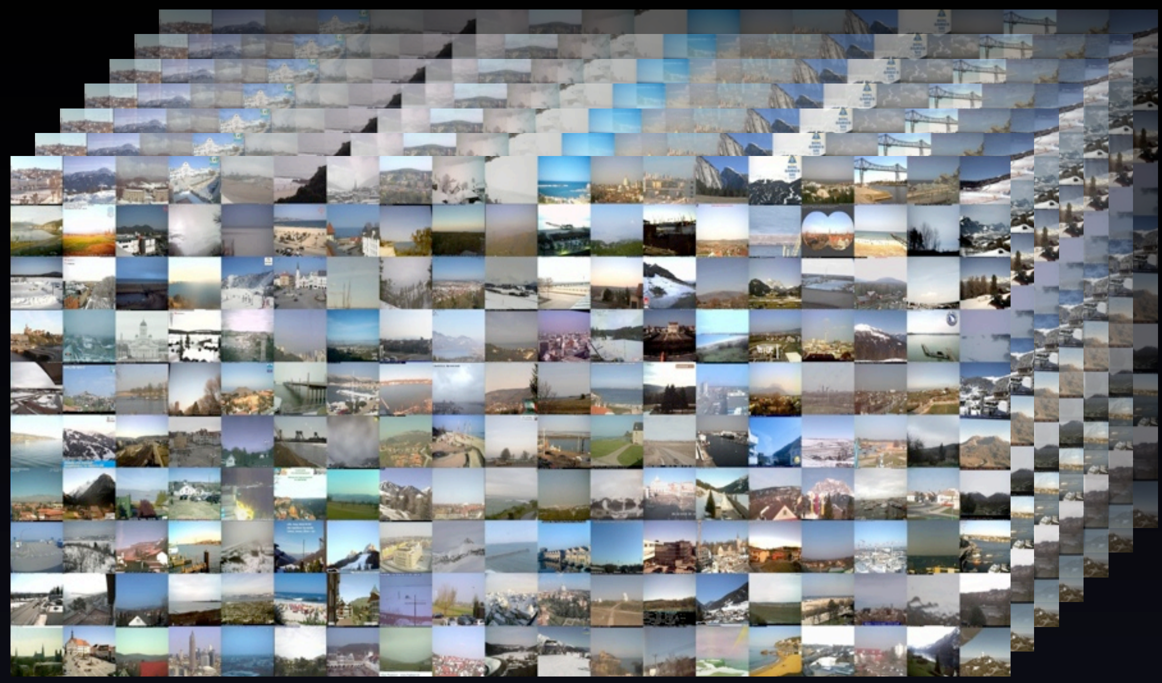

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

To conclude, we’ve presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it’s a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

Database of image sequences

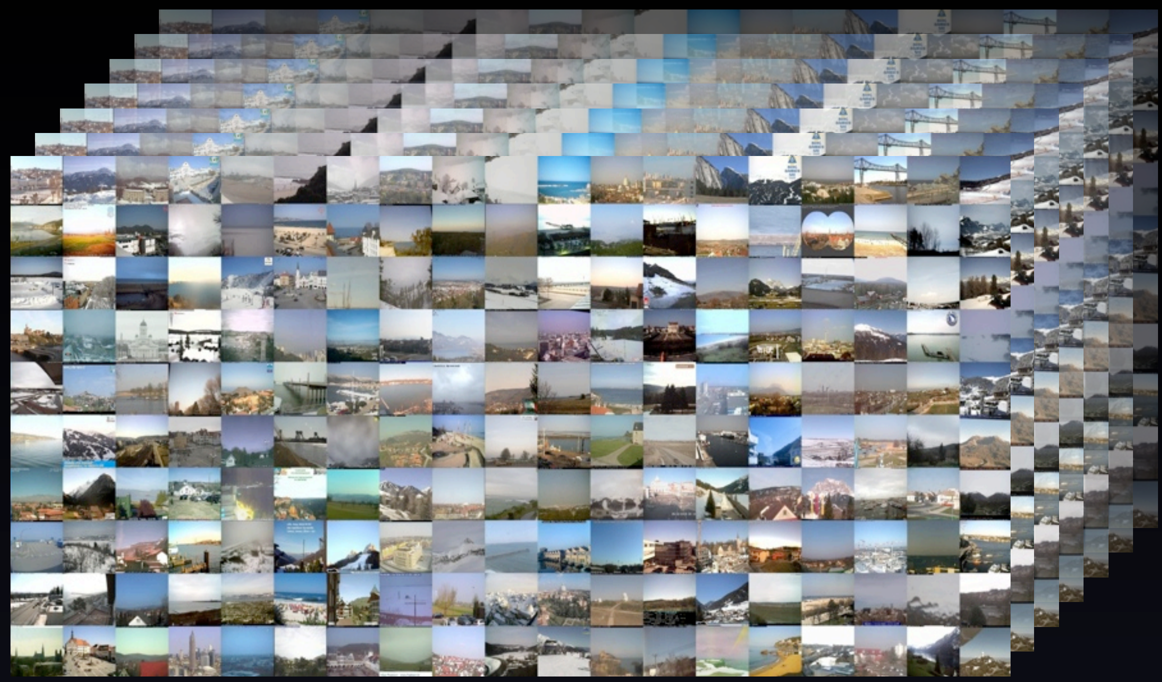

Illumination-aware scenes

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

To conclude, we've presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it's a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

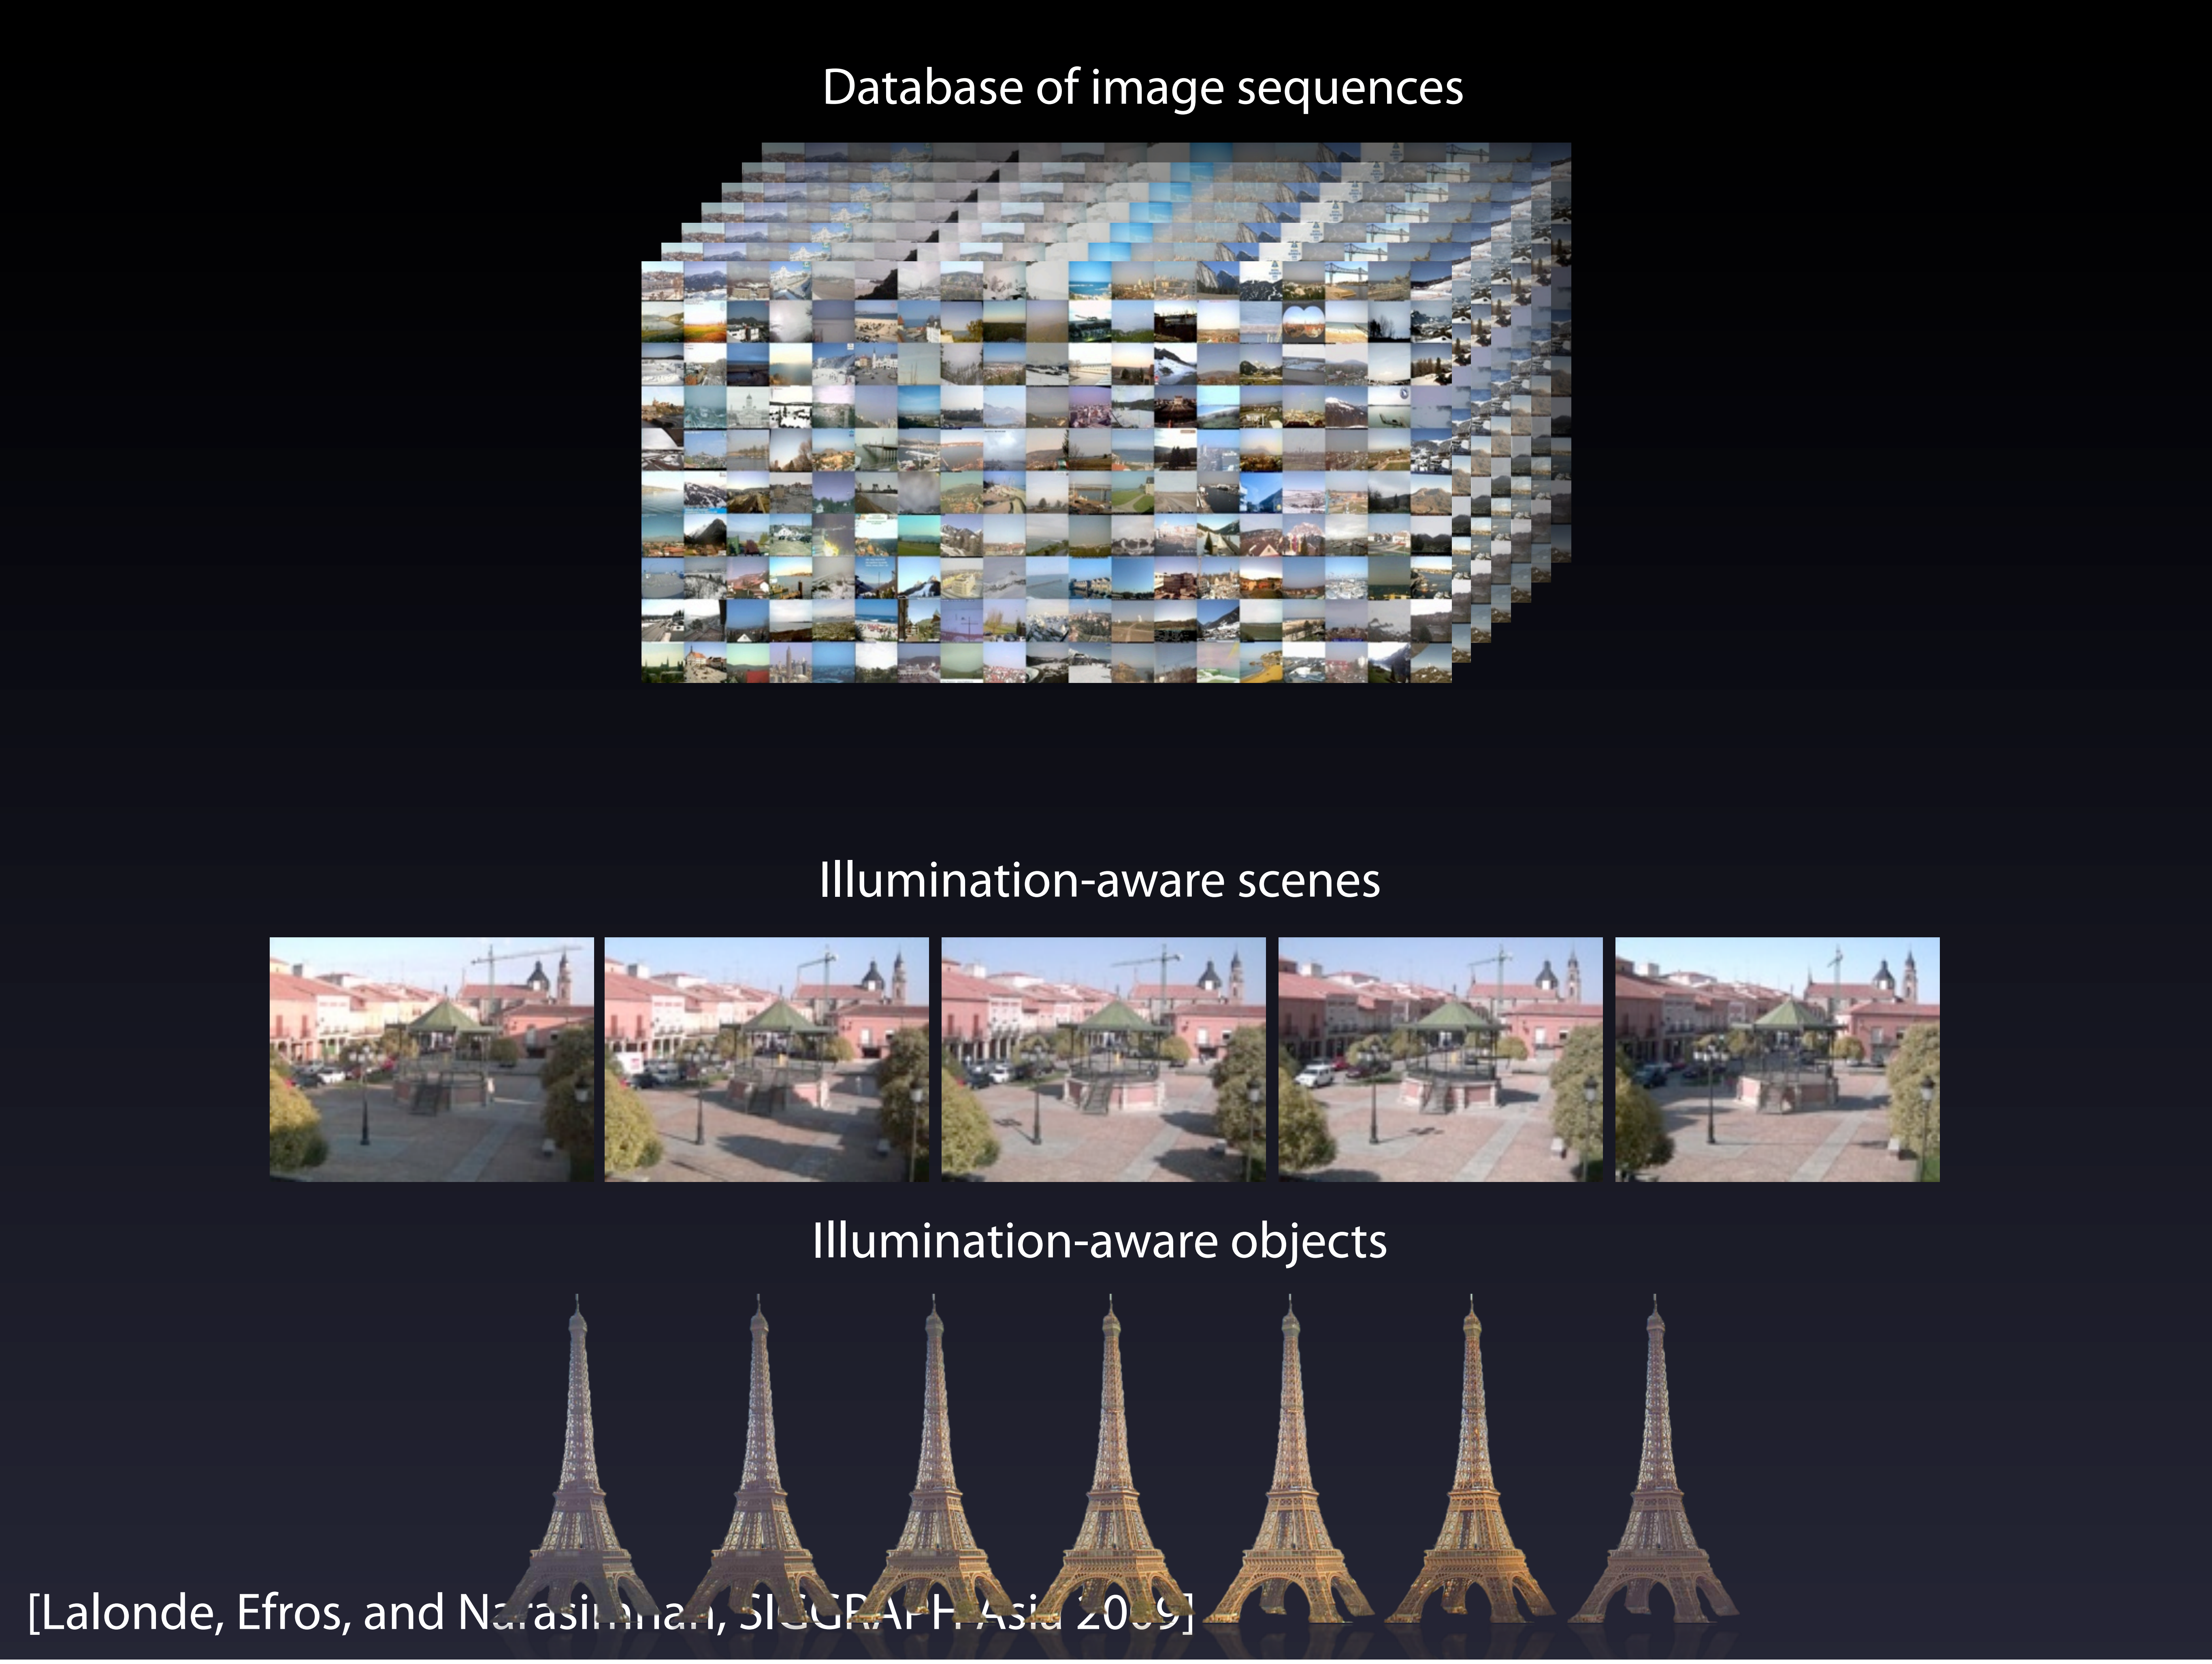

To conclude, we've presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it's a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

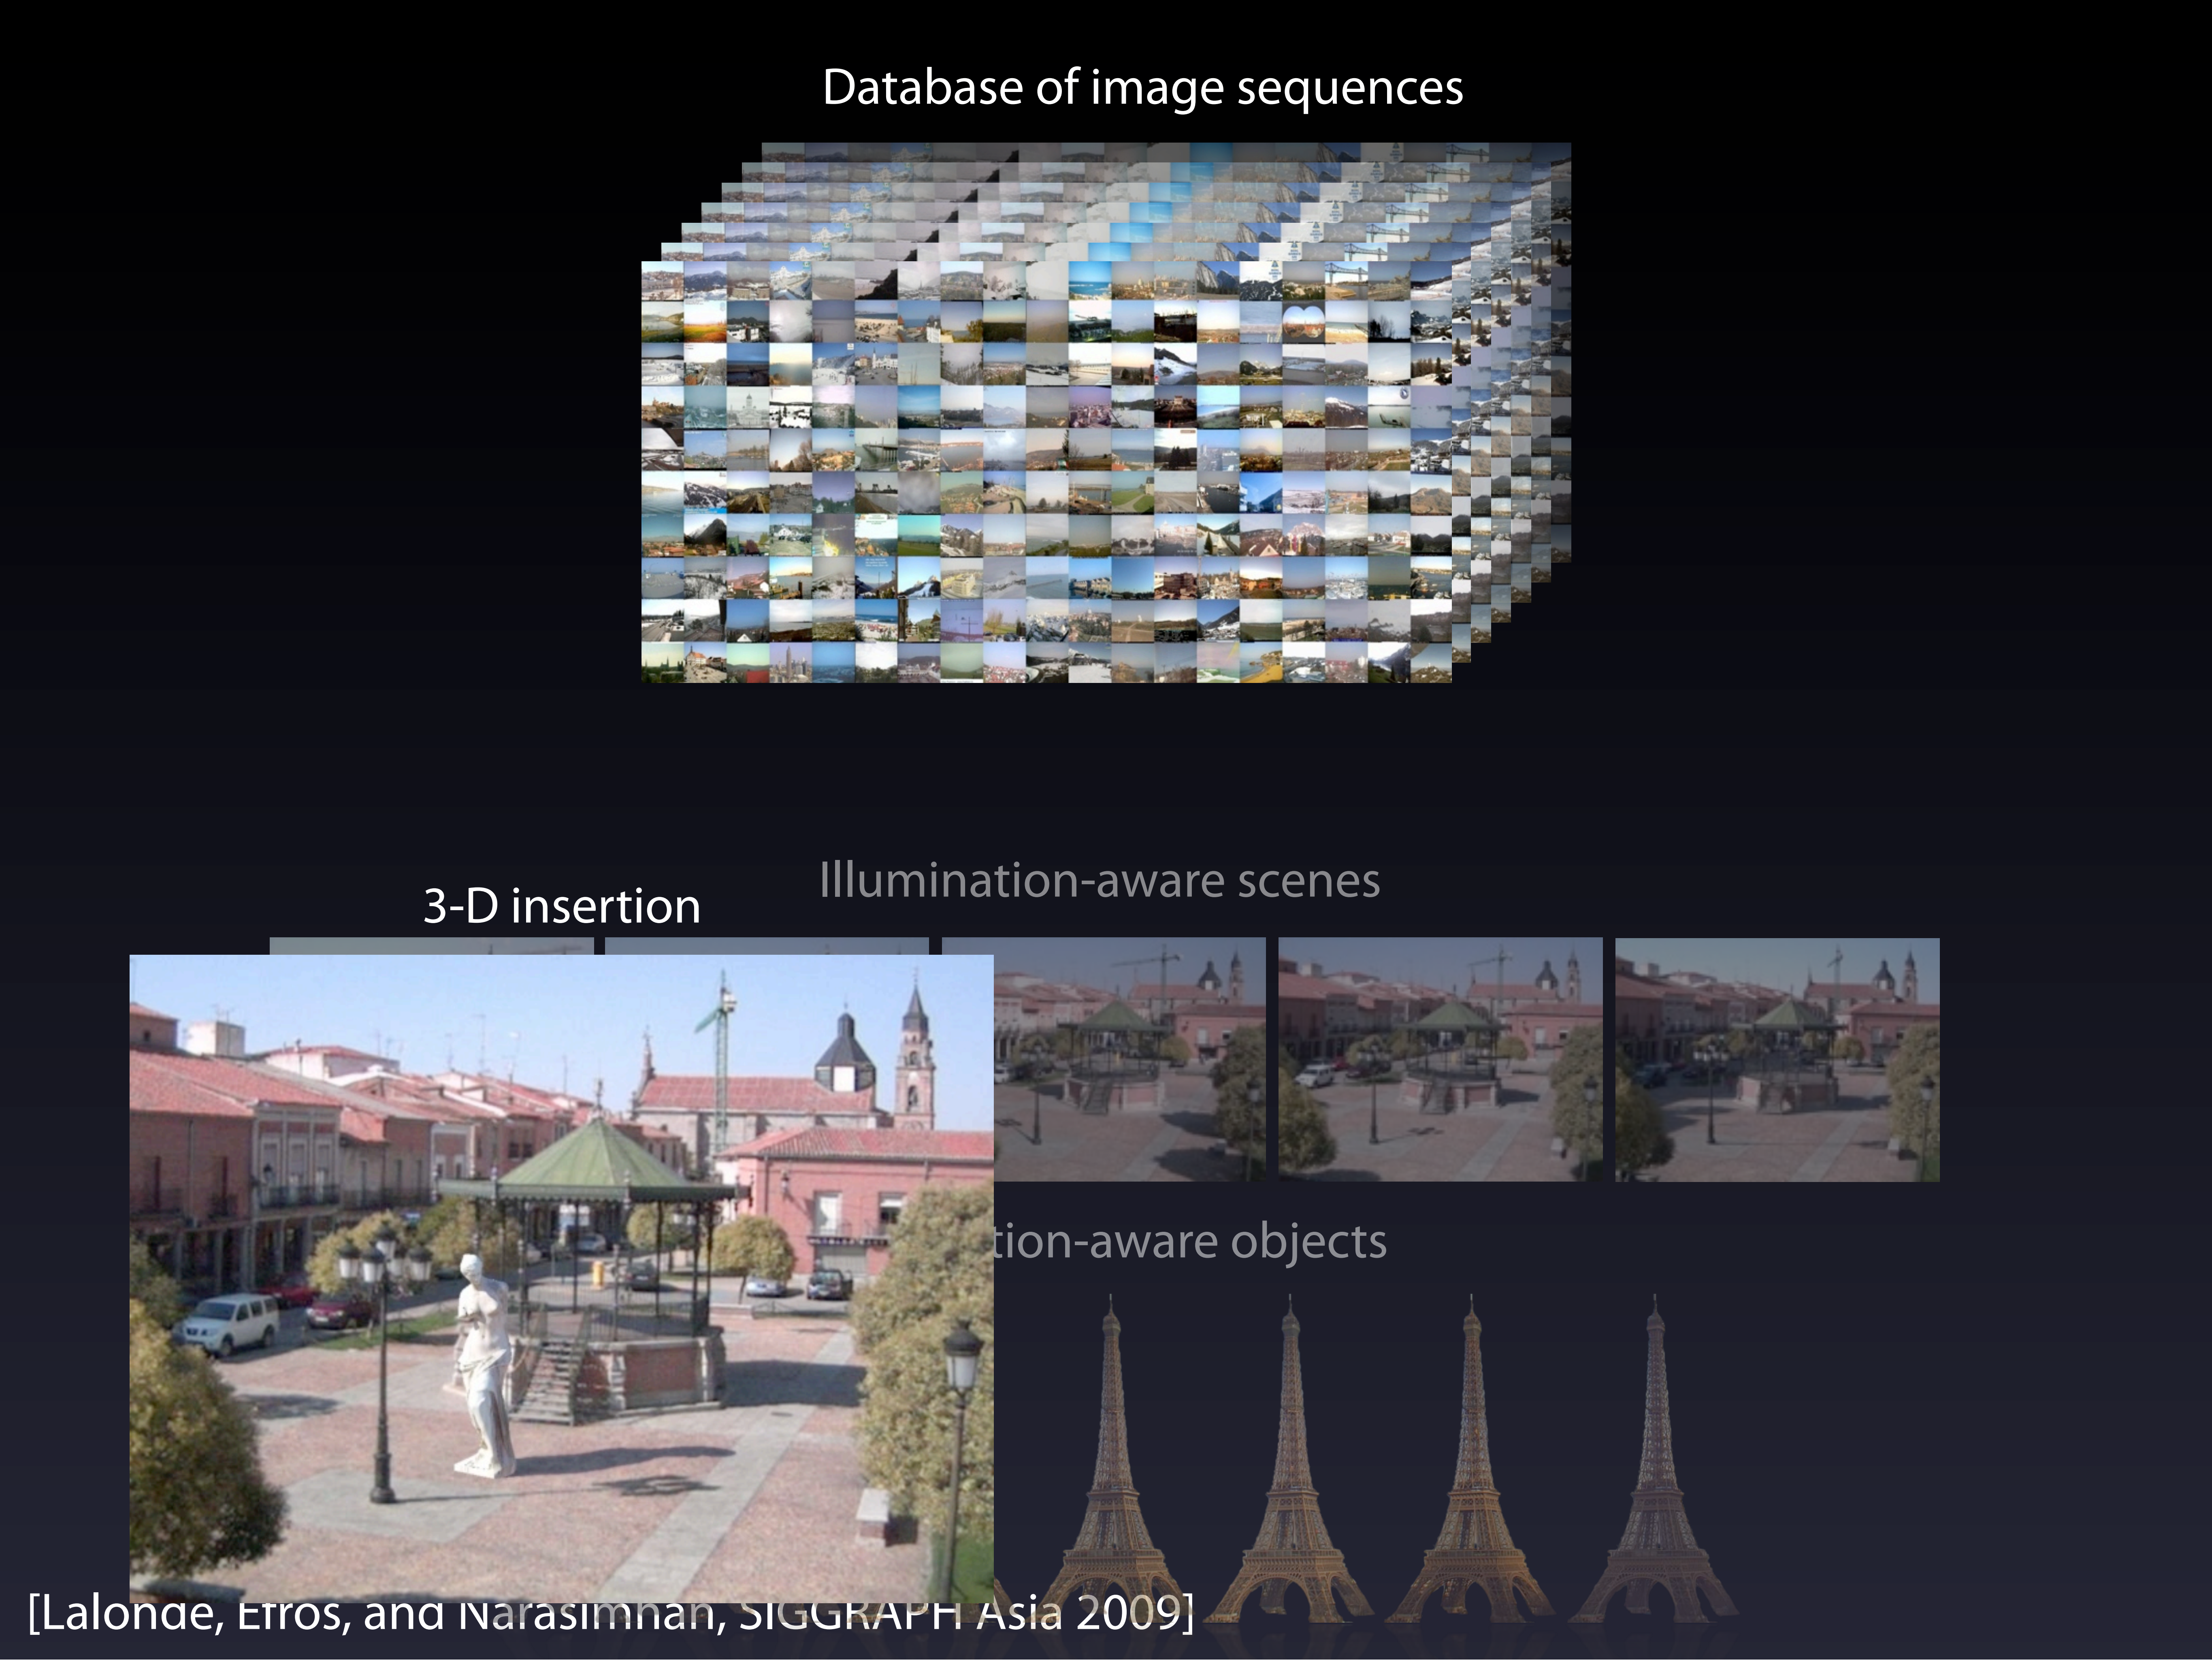

To conclude, we've presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it's a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

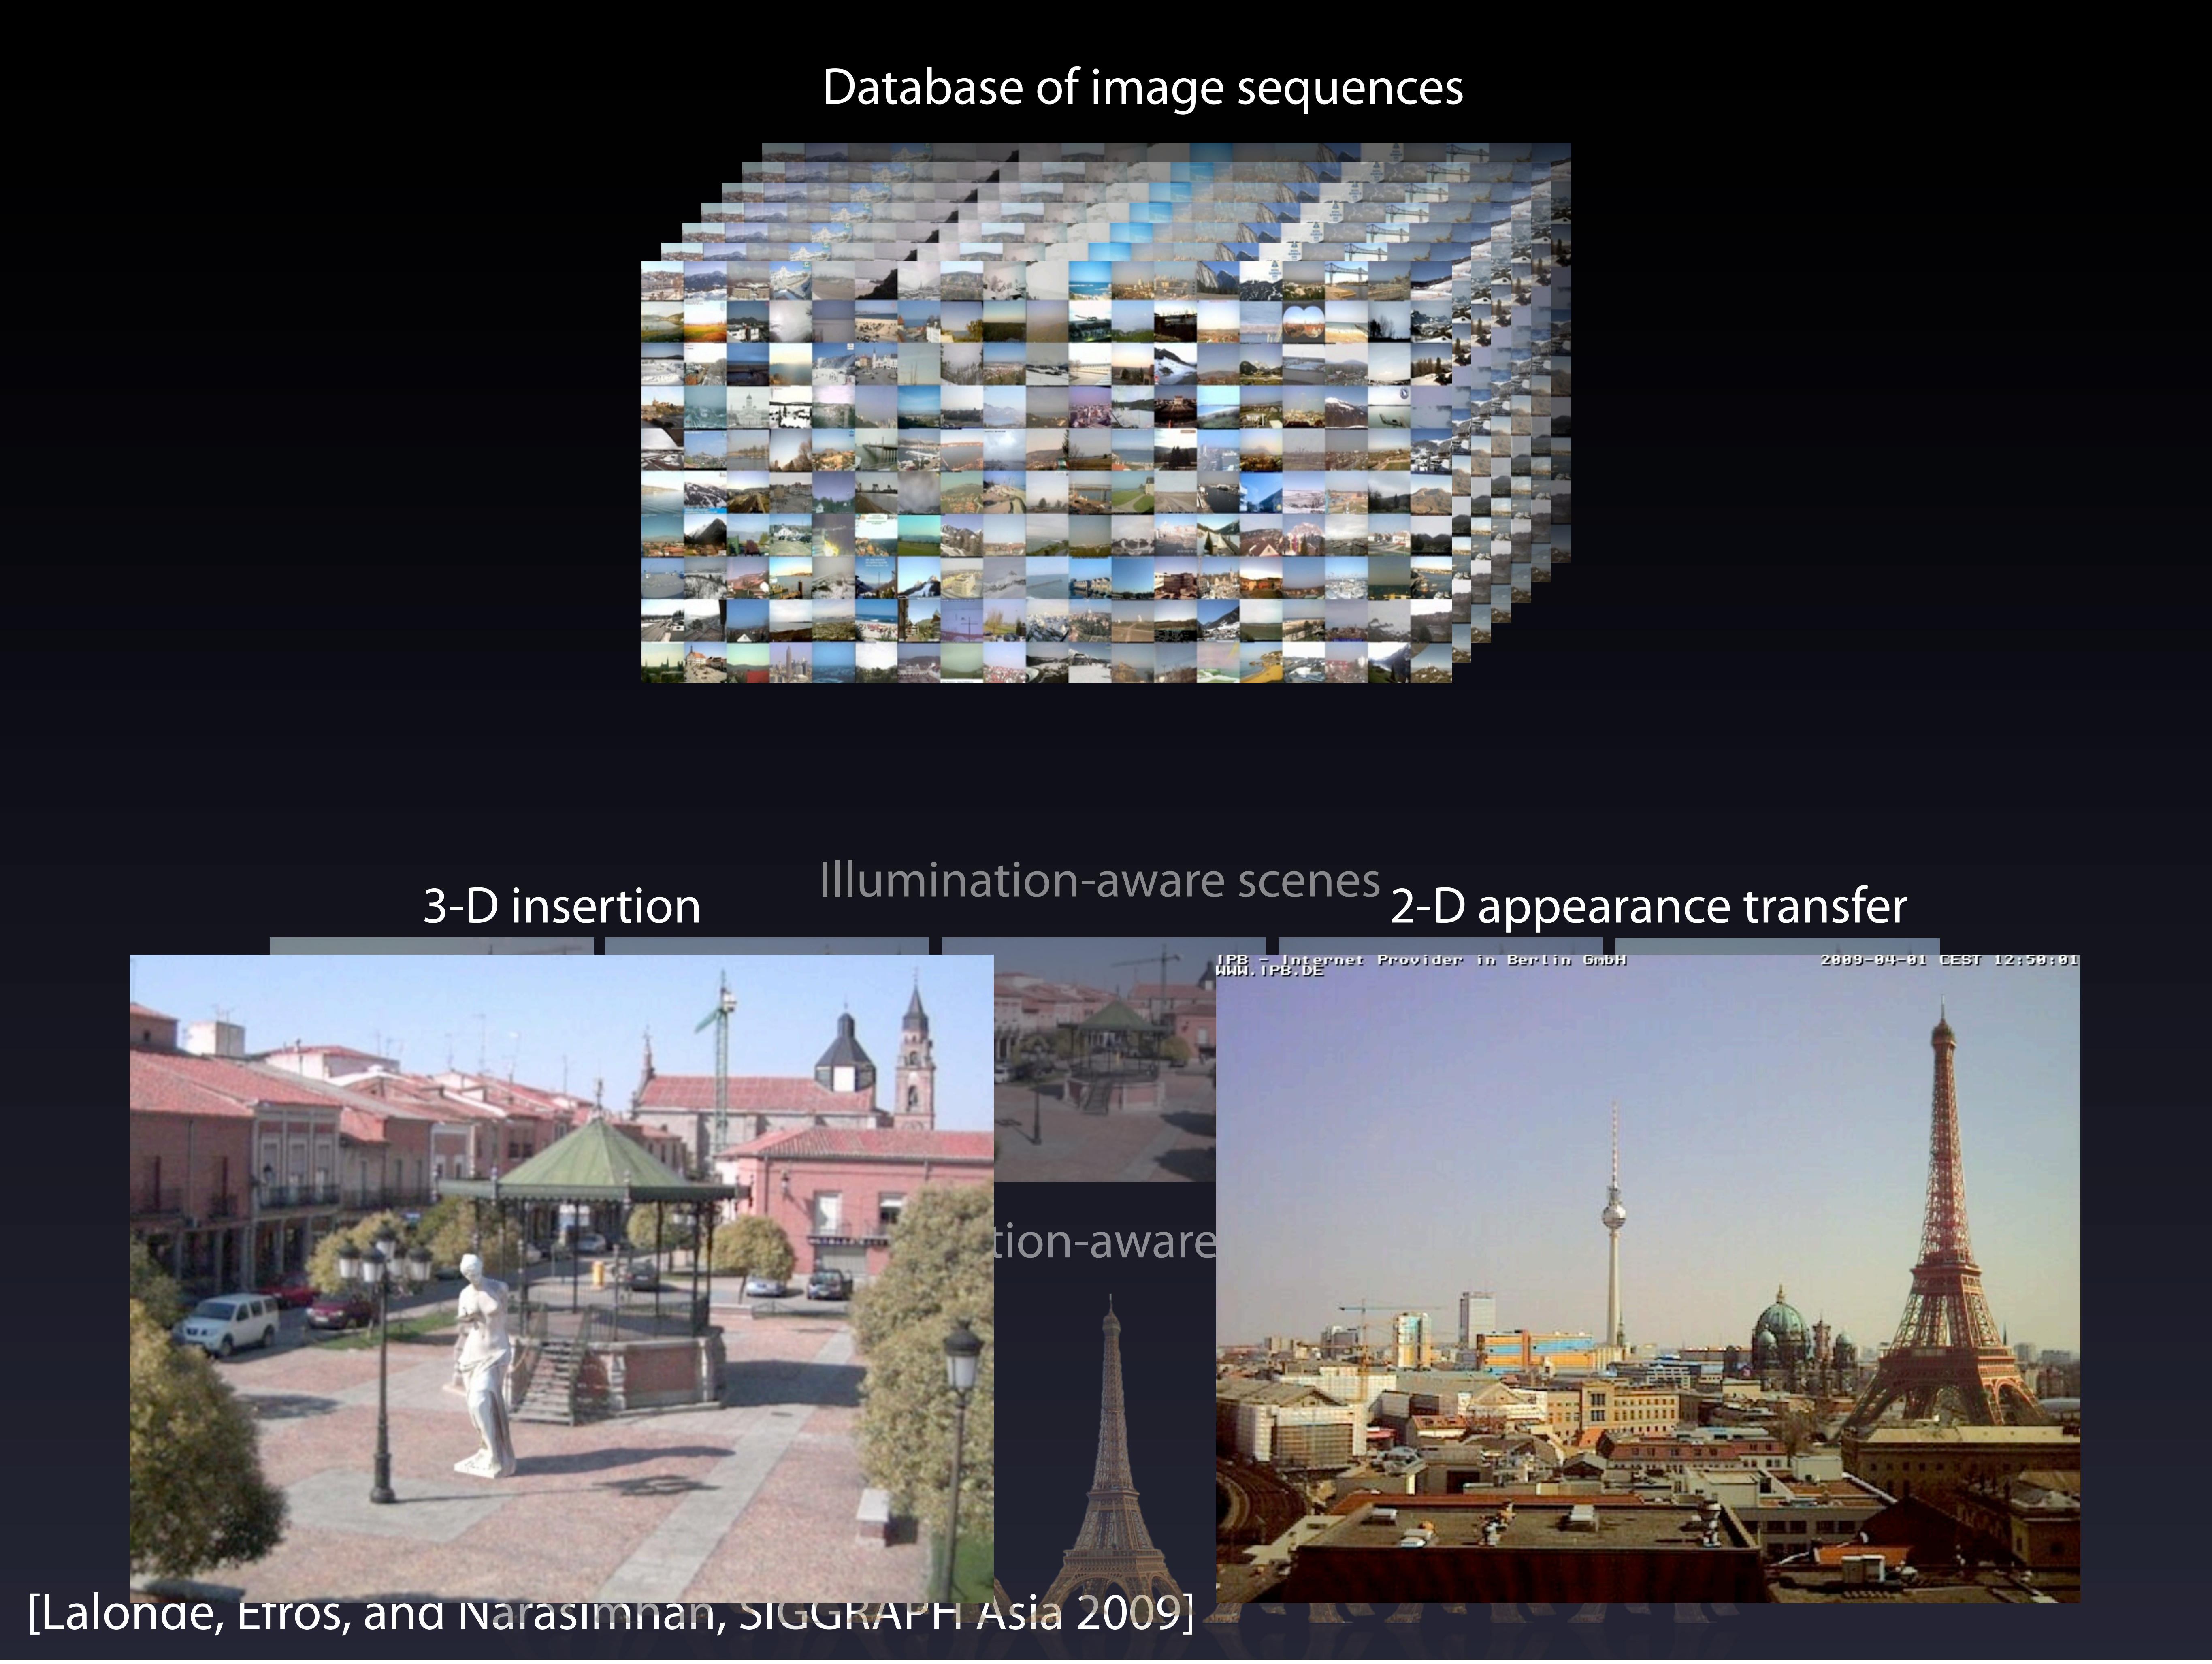

To conclude, we've presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it's a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

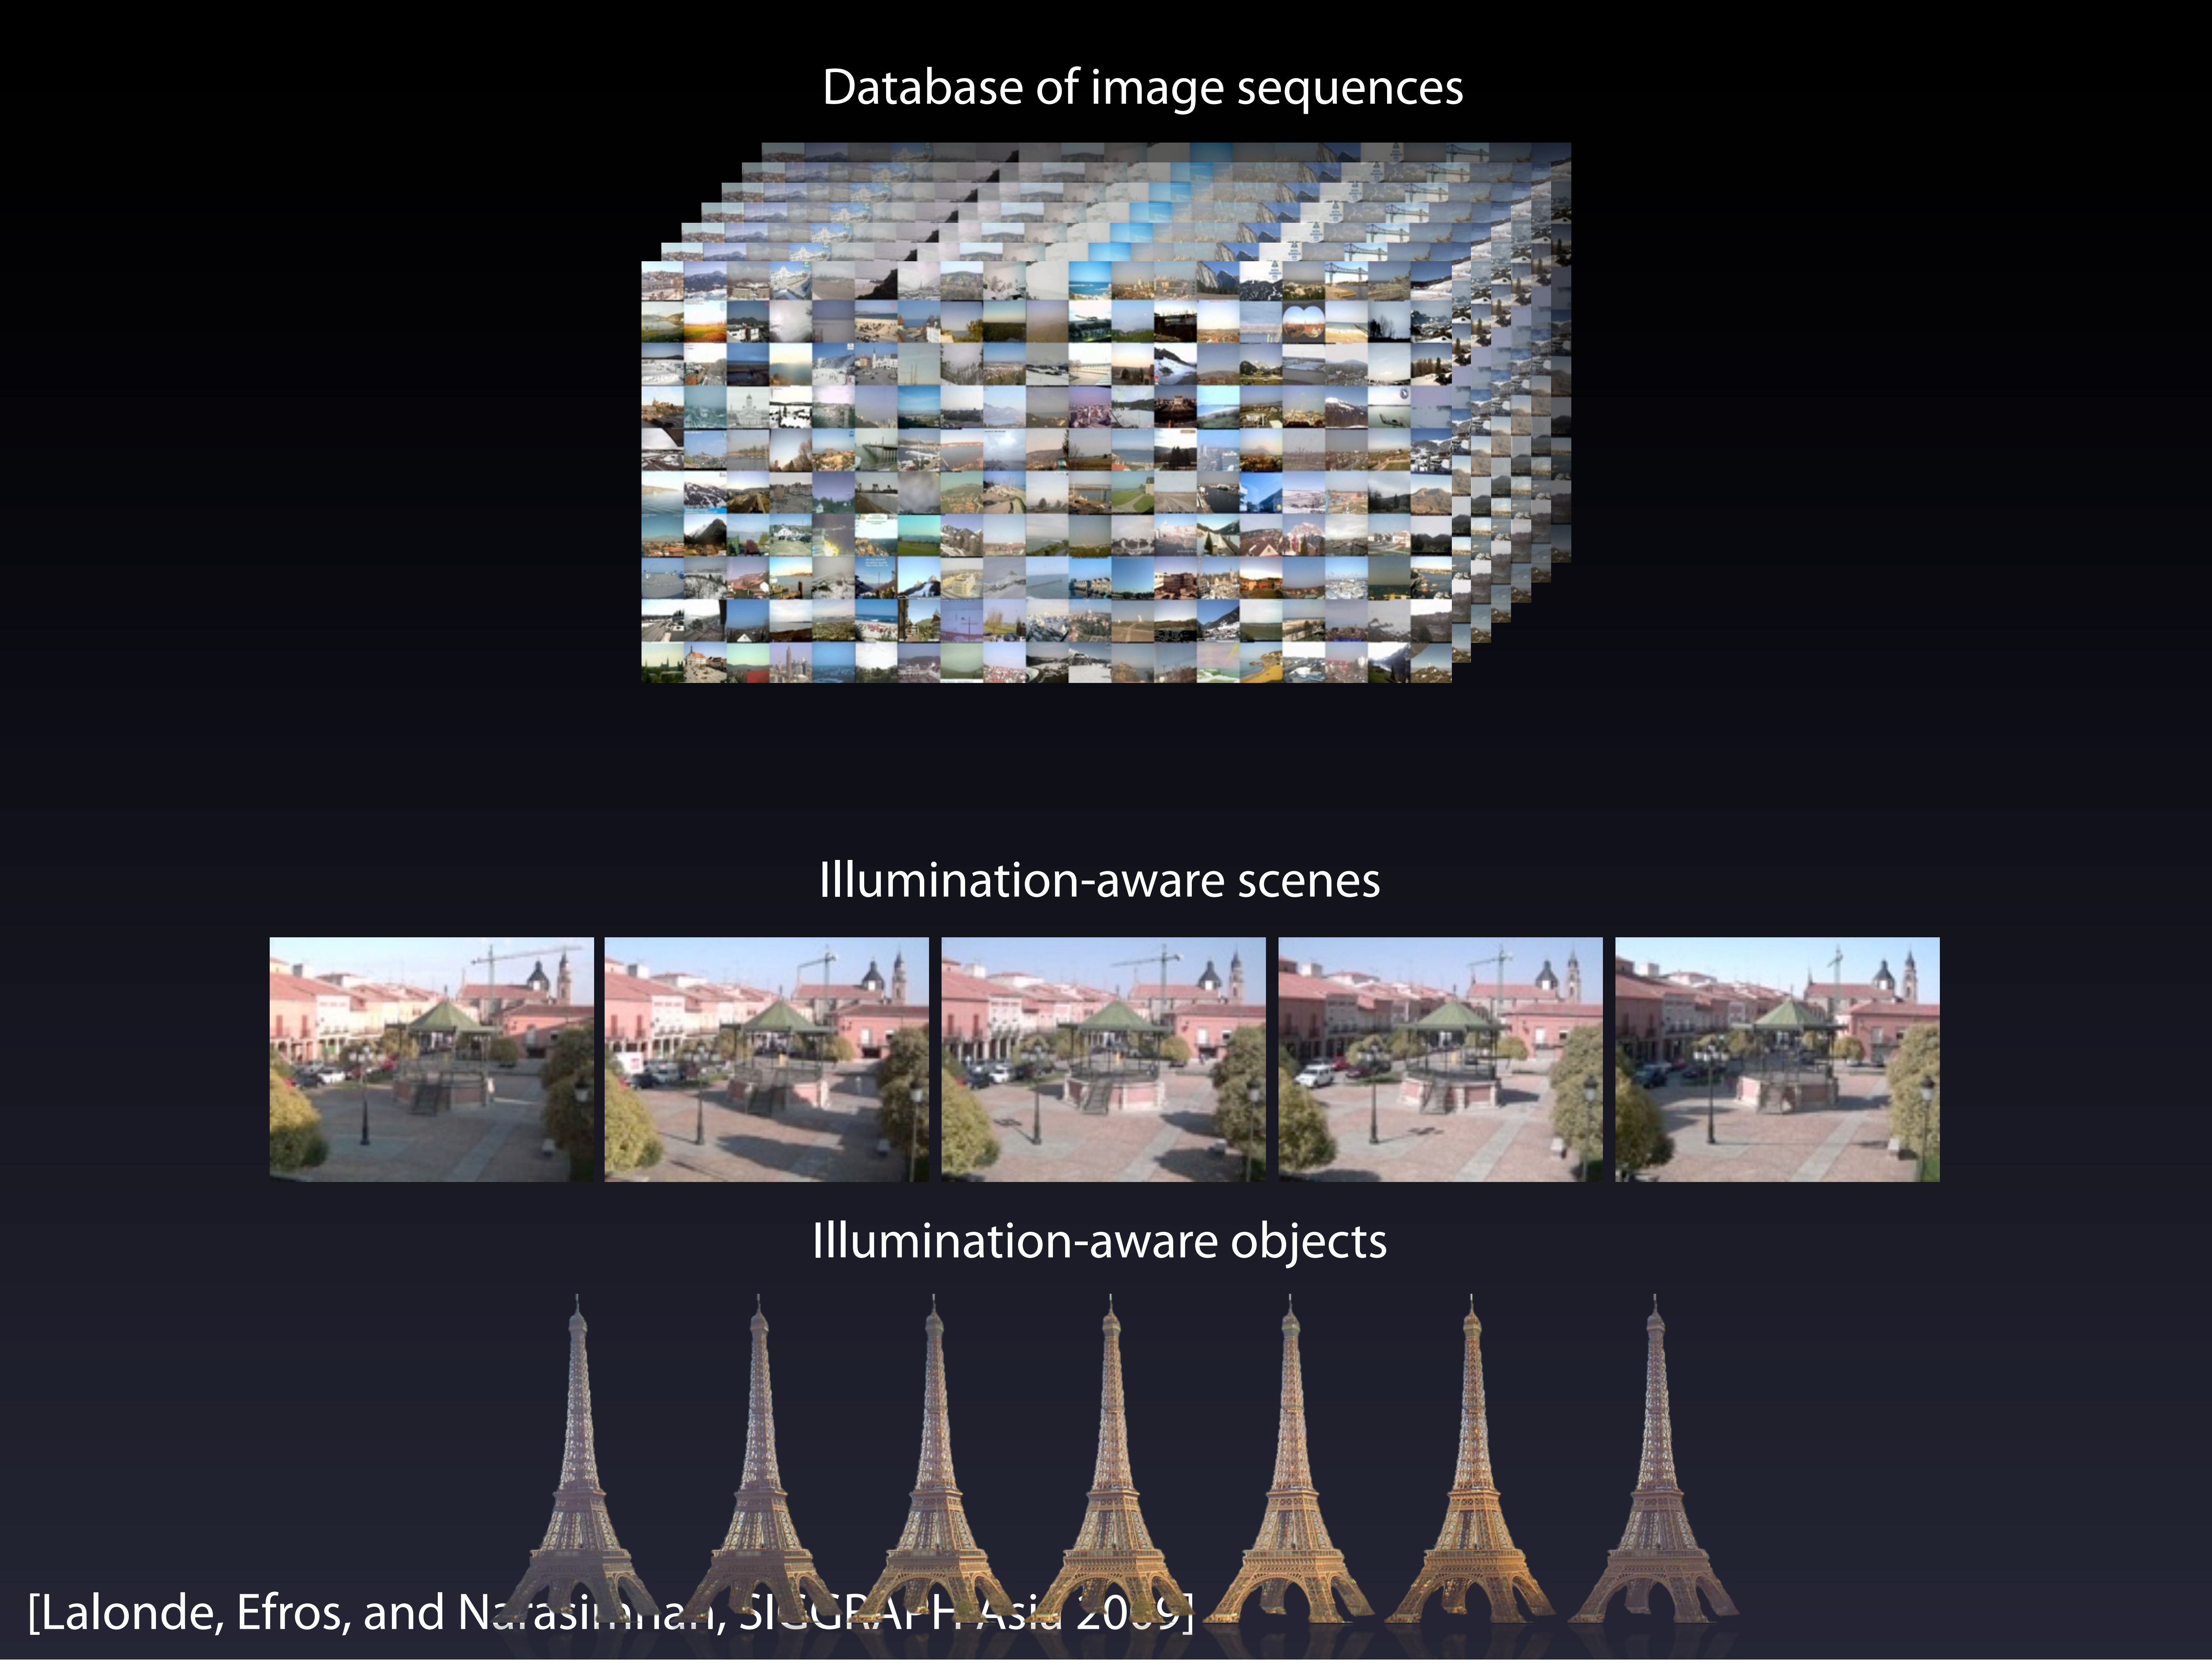

To conclude, we've presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it's a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

To conclude, we’ve presented a new database of webcams collected from all over the world, and shown how this new source of data is much more than a set of images, it’s a set of scenes and objects that can be made aware of their illumination conditions.
We showed how we can use them in 2 main applications: 3-D object insertion, which required us to automatically create a virtual environment map; and 2-D object transfer, which can be performed by matching illumination parameters across different cameras.
We hope that this will spur interest in the research community to go beyond these applications, and tap into this new source of data that still has much to reveal.

Extra slides

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Object transfer – sky

Source

Destination

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Finally, the sky itself can be seen as an “object” which we wish to transfer. In this case, we copy the sky, but we want to know how would the scene look under such a sky.

Object transfer – sky

Source

Our result

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Finally, the sky itself can be seen as an “object” which we wish to transfer. In this case, we copy the sky, but we want to know how would the scene look under such a sky.

Object transfer – sky

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Here's another example. Notice how the position of the sun is consistent with the reflection on the water, and how the scene becomes darker at sunset.

Object transfer – sky

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Here's another example. Notice how the position of the sun is consistent with the reflection on the water, and how the scene becomes darker at sunset.

Light stage objects

[Lalonde, Efros, and Narasimhan, SIGGRAPH Asia 2009]

Alternatively, we can also an object captured in a light stage, as introduced by Debevec et al in 2002.

Data source: webcams.travel

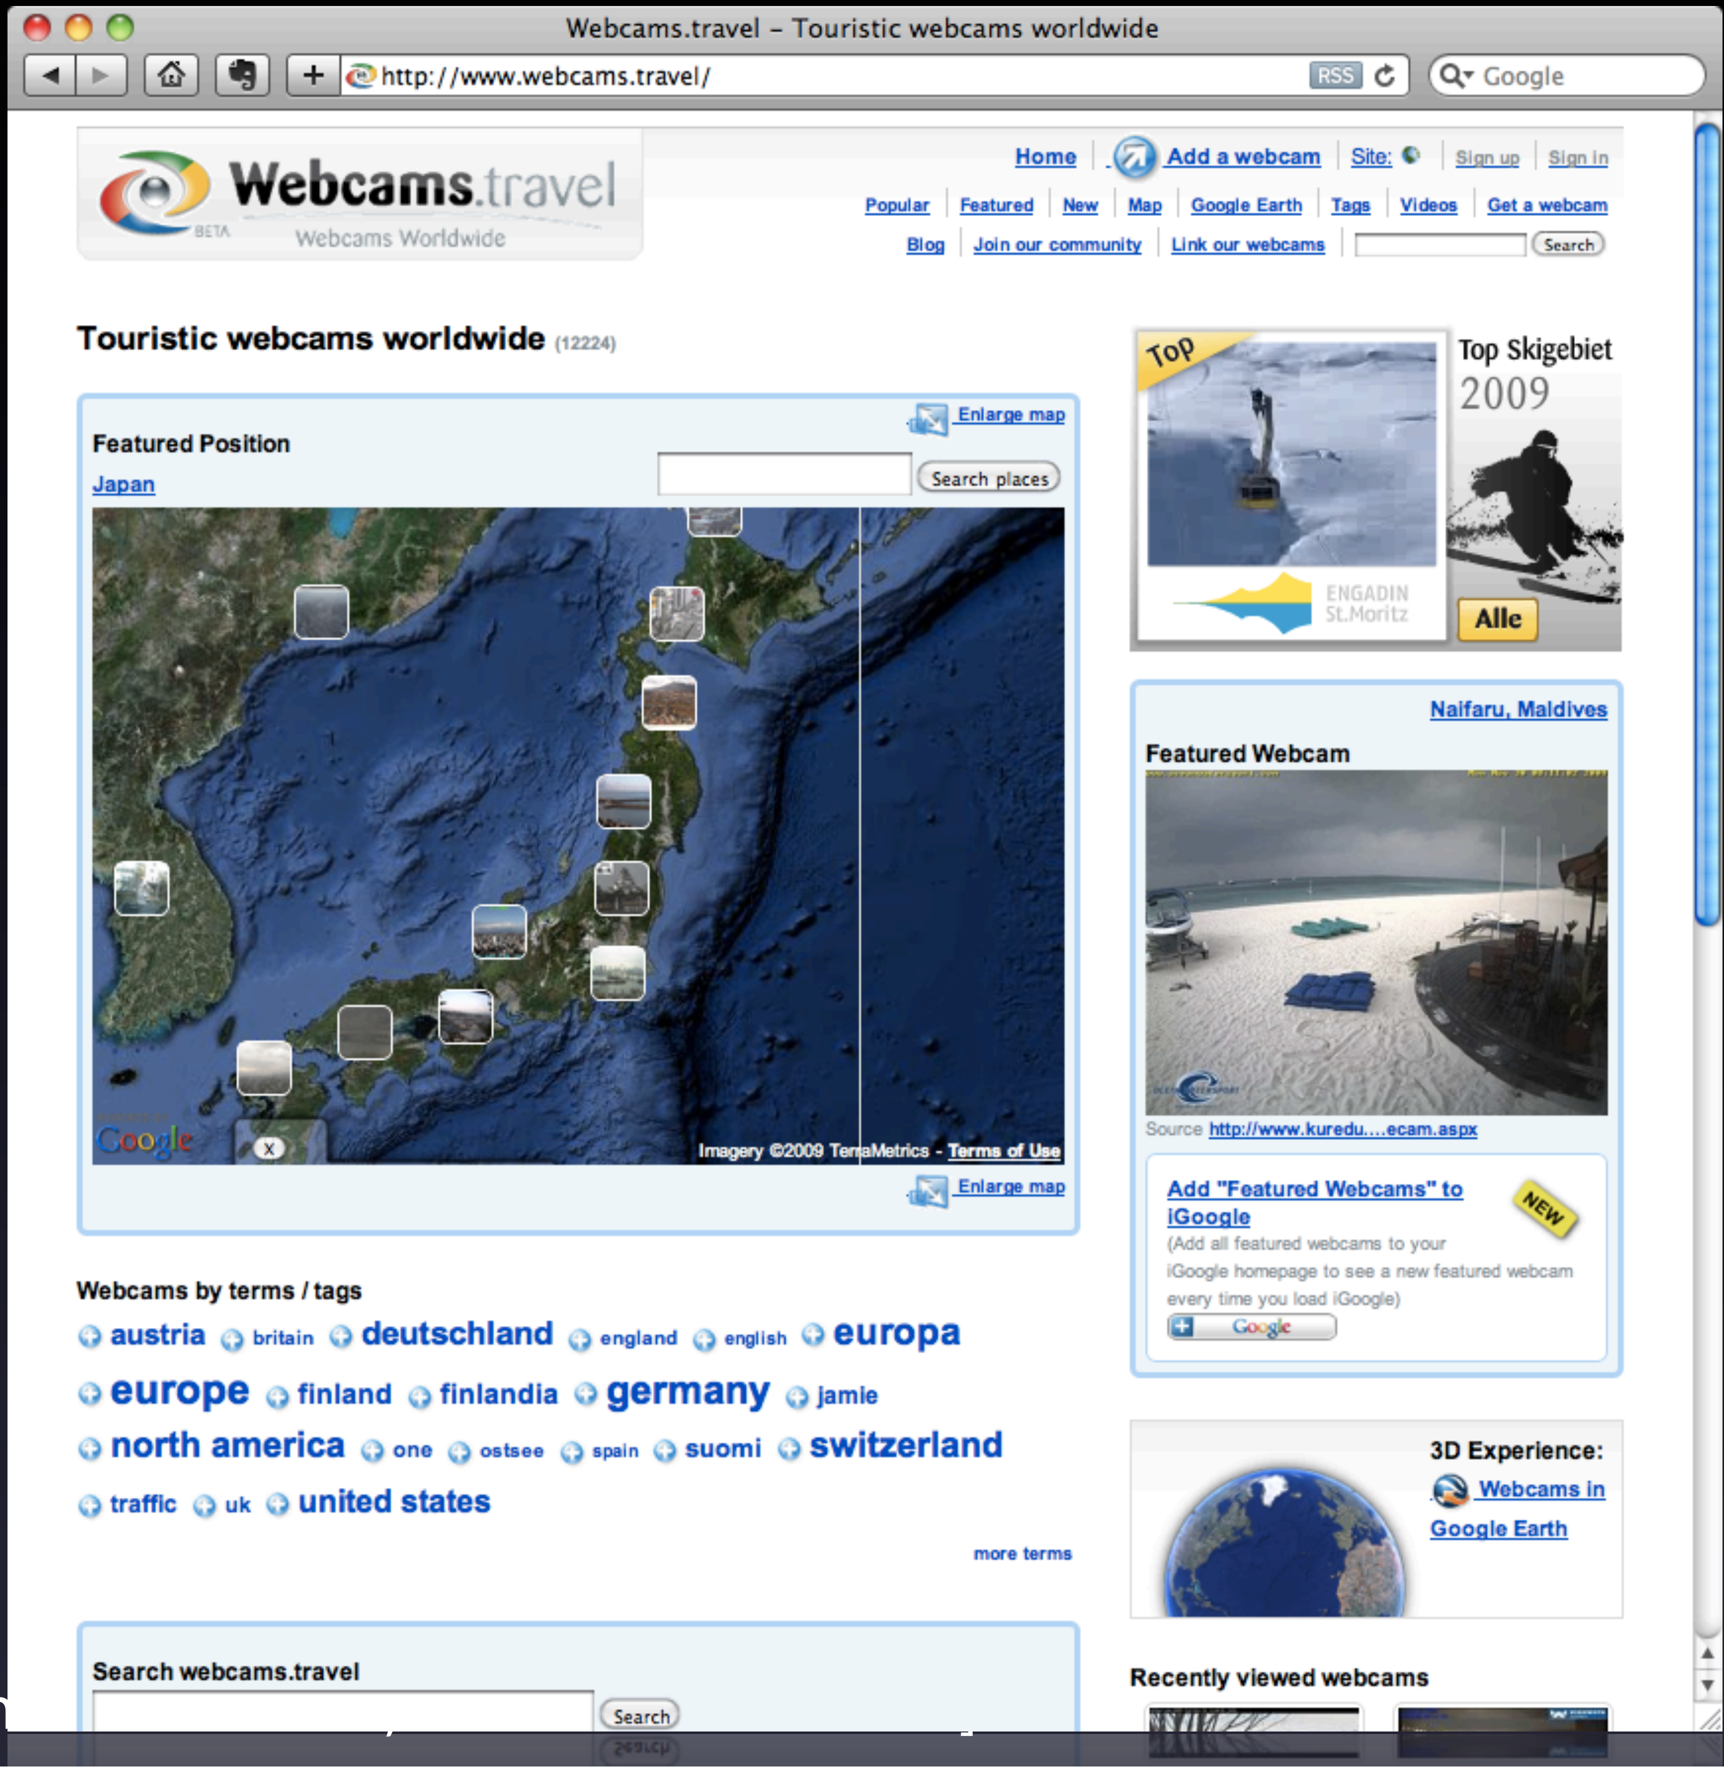

[Lalonde, Efros, an