Associate Editor
Reviewers agree that the paper makes an interesting and novel contribution. However, there are still a few issues, which especially reviewer 5 has pointed out. I therefore recommend accept with minor revisions. In particular, I agree with the reviewer that the emphasis on direct illumination should be toned down, or comparisons with other methods have to be provided (e.g., light cuts, row-column sampling, etc.). Clarity of writing also needs to be improved, see reviewers' recommendations. In the results section, comparisons need to be improved. Error values should be provided, as well as color coded difference images as in figure 11. The artefacts in the video, as mentioned by one reviewer, should be discussed (where do they come from?).

Reviewer: 4
Recommendation: Probably Accept after Minor Revisions

After the major revision, the paper contains more discussion about the proposed method and most of my questions are answered.
The paper looks acceptable to me now with some minor corrections:

Villa Arpel scene results:
You mention that 4000 shadow rays are used for the artificial light sources. Please mention the total number of artificial light sources that this scene contains. I think it is much lower than 4000, so please explain why so many rays are necessary.

The Arpel scene contains 11 area light sources represented by 760 polygons. Each area light source is sampled using a certain number of shadow rays for evaluating the direct contribution. The total number of shadow rays is equal to 4000. For the Villa Arpel scene, our experiments showed that 4000 shadow rays gives good quality results.

Figure 4: You mention that the circular records require the same computation time as the new, adaptive records. This is difficult to understand since the circular records require the determination of a mean distance. This is also required for the adaptive records, and additionally an iterative procedure is performed for each of the eight axes (Figure 9). So please explain why there is no overhead.

The iterative procedure needed fot footprint adjustment, depending on the irradiance value, has not been performed for figure 5, which explains the fact that the computations of the circular and adaptive records take the same time. A mention about that was added in the figure 5 caption.

Starting from Figure 9, the ordering of the figures is messed up.

The order of the figures is respected in the new version of the paper.

Typos:

These errors have been taken into account and corrected in the new version of the paper.

Reviewer: 1
Recommendation: Probably Accept after Minor Revisions

Thank you for the experiments on varying rho and curved surfaces.

The primary point which the article makes, that using the new definition of neighbourhoods helps to reduce

the number of records stored drastically particularly in the presence of high-frequency illumination effects, is convincingly presented.

However, looking at the quality of the shadow edges in submitted video, it is not obvious that this gain in the number of records is necessarily good. Particularly at times: 0:44 (chair shadow), 0:53 (the vertical edges of the glass window boundaries and near the base of the column closest to the camera and also on the deck umbrella). Is this due to some bug ?

I think that, before this work is publishable, the source of these artifacts needs to be confirmed and it has to be verified that they are not inherent to the reduction of samples by the method.

The artifacts visible at time 0:53 are due to bugs in our lighting simulation software and are not inherent in the reduction of the number of records (see my previous comments to reviewers). Concerning the poor quality of the shadow edges (at time 0:44 for chair shadow), this issue is due to a permissive rho parameter. A more restrictive rho parameter would give better results. However, this would increase the number of records as well as the computation time. Let us recall that our software is dedicated to physics-based lighting simulation where the objective is to compute physical quantities, such as irradiance and radiance, needed by lighting engineers. Visible artifacts are then tolerated.

Reviewer: 5
Recommendation: Resubmit with Major Revisions

== Overall judgement: Average ==
This paper presents improvements on the Irradiance Caching technique, namely adaptive non-circular footprint of each record and a new gradient computation method based on equivalent light sources. The results shown in this paper demonstrate that the proposed method can use a fewer records to achieve perceptually similar quality of images (Figure 5, 10, and 14).

Overall, I think changing the shape of the footprint to something other than a circle is a clever improvement and certainly novel in terms of the application to Irradiance Caching. Although the current technique seems not applicable to non-diffuse surfaces, its contribution sure has immediate impact in some practical usage. However, I think the presentation and demonstration of the results are slightly below acceptable level for publication in its current form. There are some important information that is missing for judging results as I will mention in the following. The paper needs a few more major revisions before its acceptance in order to accurately convey its contributions.

== Originality/novelty: Good ==
The non-circular footprint in Irradiance Caching is definitely novel as far as I am aware of. However, as one of the reviewers pointed out, I think it is very closely related to 'Improved Illumination Estimation for Photon Maps in Architectural Scenes' by Tobler and Maierhofer. I agree that their paper is for photon density estimation, but the kernel shape is essentially the same and the motivation of using this shape seems to be the same as well (capturing changes in illumination around edges and corners more accurately). It is definitely worth mentioning in my opinion.

A mention to the paper, entitled 'Improved Illumination Estimation for Photon Maps in Architectural Scenes', has been added in section.

The gradient computation seems to be novel as well. I would however appreciate an explanation on why the original method by Ward is not appropriate in combination with the proposed non-circular footprint. As far as I understand, this original method by Ward does not assume any smoothness of illumination, thus including rapidly changing illumination should not matter. I would also think that changing the shape of footprint should not affect gradient computation, because gradient of irradiance should be independent of the

footprint.

Yes we could use Ward gradients. As these gradients are of the first order, the irradiance changes are considered as smooth. For high irradiance variations, we need second order gradients, which is the case in our method as direct contributions are also stored in the records. This ~~has~~ pointed out in the introduction of the section 5.

== Importance: Average ==
Since Irradiance Caching is widely used in movie industry, technical improvements proposed in this paper may be quite useful in many practical applications. However, based on results in this paper, the benefit seems to be rather limited to scenes with relatively smooth direct illumination that also dominates computation time or scenes with non-smooth indirect illumination where existing methods generate too many records. The method is also restricted to Lambertian surfaces. The authors mentioned glossy BRDFs can be handled in the second pass by tracing rays similar to final gathering, but Radiance Caching might work better in such cases.

== Technical soundness: Average ==
Using a different interpolation kernel is technically sound. In particular, since the one used in this paper is similar to the one proposed by Tobler and Maierhofer in a different context but for similar purpose, the proposed interpolation kernel looks like a natural but nontrivial improvement over the original Irradiance Caching technique. Unfortunately, the part that describes adaptation to illumination changes is rather terse (Section 6). I would expect more detailed explanation of this point since it seems to be one of the important changes from the original Irradiance Caching technique.

We agree with you. We have clarified this point accordingly.

The new gradient computation is sound too, but it lacks some technical solidness. For example, you use the term 'point lights' for indirectly bounced illumination from diffuse surfaces. They are *not* point light sources at all since they have directionality due to normal directions. The authors might be aware of this (is that why you have differences in Equation 13 and 14?), but at least the current writing is very confusing on this point.

This point has been clarified ~~at~~ section “Background”. This section has been previously named “Overview”.

Another point is the visibility term in Equation 12. It is unclear, if we need to evaluate this term by additionally tracing rays toward some points along axes of the zones, or we just assume it is one.

This point has been clarified ~~at~~ sections “Background” and “Computing Irradiance along axes”.

In other words, it is unclear how equations interplay in the final algorithm. Derivations of Equation 13 and 14 need to be elaborated (how you get to those equations? what is the assumption if any?).

Equations 13 and 14 allow to provide an approximation of the intensity of an equivalent point light source.

It is also unclear if you are using gradient, which is first order derivatives, or higher order derivatives as well. Equation 15 seems to indicate that we use second order derivatives, which are not gradients.

To account for high frequency gradients, we use second order gradients.

As I mentioned, it is also unclear why the original method by Ward fails. Ward seemed not assuming low frequency signal in gradient computation as far as I understand as opposed to what's mentioned in this submission.

If we use a first order gradient, we implicitly assume that the irradiance changes are smooth, which is the case of Ward's gradients.

== Clarity of writing and presentation: Poor ==
This is the part that largely affected my decision. This is also somewhat related to technical soundness above since unclear writing and presentation makes technical soundness rather dubious in general.

First of all, I do not think the method should be presented as a way to incorporate direct illumination in Irradiance Caching technique. I essentially agree with Reviewer 2's comment included in the supplementary material. The proposed method can handle direct illumination better than existing methods simply because it improves placement of records and the interpolation stage. The improvements have virtually nothing to do with direct illumination, or the method does not exploit anything related to direct illumination. For example, if direct illumination is smooth enough (e.g., hemispherical constant illumination), existing methods will work well even with direct illumination. If the authors still consider the method has something special for handling direct illumination, direct comparisons with other existing work, such as Bala et al. is must, not just in text. In either case, the abstract, related work section, and result section need major changes accordingly.

Let us recall that the goal of our method is to cache effects due to high-frequency illumination. Our method can be used only for caching diffuse indirect illumination with high frequency changes. In this case, ~~the~~ our adaptive method decreases significantly the number of record, especially near the corners (thanks to the adaptive zones) and allows the caching of effects due to strong indirect equivalent light sources. Storing direct contributions can be seen as a extreme case of indirect high frequency illumination.

In general, comparisons in the paper lack of objective measure. For example, I understand that Figure 5 shows that the proposed method results in a fewer number of records with the same parameter value that needs less computation time to fill the records. However, it is wrong to claim that this is a positive result because we do not know if the result of the proposed method has the same error as the result of the compared existing method. If the error has increased, the result is not worth mentioning since having a fewer points naturally leads to larger error. In other words, generating fewer records does not mean improvement, unless computed error (either perceptual or purely numerical) is equal or even smaller. Figure 10, 12, 13, 14, and 15 all should provide computed error values so that we can verify that the proposed method did not increase error by having a fewer records. Figure 11 is good, but the similar results should be provided for Figure 10 so that we can see what kind of error we get from each method. Figure 8 should provide rendering time, as the proposed method might be better at the cost of additional computation. Objective comparison is lacking throughout the paper in my opinion. The text often says results are visually equal which is subjective and sometimes subtle.

We agree with you. We could add more results validating our method but that would increase significantly the size of our paper which is already of 13 pages. Our method is already used for many project of lighting simulations.

I list more detailed questions and comments in the following.

= Abstract =
* '(2) to store both direct and indirect'
I disagree with claiming this as contribution as I mentioned above.

We replaced this sentence by “to efficiently render the high frequency illumination changes.”

* 'With this technique, the record density will no longer be uniform.'
Remove this sentence as existing methods do produce non-uniform record density.

This sentence was changed.

* ', especially animation rendering'
Mention that it applies only for walk-through animation where illumination and positions and shape of objects do not change.

Changed in the new paper version.

* 'wide surfaces exposed to daylight'
Since Irradiance Caching is a view dependent technique, absolute size of a surface should not matter. Please explain why it matters.

We agree with you. The “wide” term was removed.

= Section 1 =
* '(glossy reflection, refraction, caustics, subsurface scattering) ... To take up this challenge, '
Maybe rewrite these as Irradiance Caching does not solve any of those.

Done.

* Multiple sentences start from 'It'. For better writing style, I would replace 'It' with the thing that you actually refer to. For example, the first `It' is probably 'The original Irradiace Caching' whereas the second 'It' is perhaps 'The algorithm' or something.

It is changed in the new version.

* '(shadow map for point light sources, Monte Carlo for area light sources)'
This is not always true and application dependent. Also, Monte Carlo is not an algorithm, so use 'Monte Carlo ray tracing' for example.

Shadow map and Monte Carlo ray tracing are just given ~~for the~~ example. It is corrected in the new version.

* 'Storing the direct irradiance ... is not possible'
It is always possible, and sometimes accurate enough when direct illumination is smooth enough. The issue is that existing methods might have large error for high frequency illumination, which has basically nothing to do with direct illumination although direct illumination tends to result in high frequency illumination.

The sentence ~~was~~ changed by 'Storing the direct irradiance ... gives poor results'

= Section 2 =
* 'spatial coherence'
Adds 'of irradiance'. There are also other instances of 'spatial coherence' in the paper.

Corrected in the new version.

* 'Unfortunately, with this adaptive...'
I failed to understand why it is 'unfortunate'. It makes sense to me to have the best interpolation in the region with the highest discontinuity.

For a set of records which participate in the interpolation of the irradiance at a point P, the record which introduces the strongest discontinuity is often the closest record to P. Unfortunately, this closest record gives the most accurate irradiance at P. We have added in the paper an example to illustrate the issue.

* 'presented in the same paper'
Specify which one as you are citing multiple papers by Krivanek et al.

The paper reference was added at this point.

* Arikan et al. seems to be unrelated to your work.
* 'less compact' -> I believe 'compact' is correct.

You are right. We removed from our paper.

* The paragraph starts from 'One of the strengths of...' is unrelated since your method is not about handling animation with Irradiance Caching.

You are right. We removed this paragraph from our paper.

* 'In this paper...'
Maybe move this sentence after the description of Radiance Caching.

Done.

= Section 3 =
* Equation 4
Use different symbol for w_R(p) as it is different from Equation 1.

w_R(p) was changed in equations 1, 2 and 3 by w^{Ward}_R(p) to avoid confusion.

* 'equivalent point sources'
They are not point light sources. Maybe the authors try to distinguish this point by saying 'point sources' instead of 'point light sources', but they are still confusing. I would suggest 'oriented point light sources' as used in the Lightcuts paper by Walter et al.

This point has been clarified ~~at~~ section "Background" as we mentioned above.

= Section 4 =
* 'it is important not to increase the density near edges and corners.'
Why? Since illumination can changes rapidly around edges and corners, it seems natural to have higher density around those regions. If you agree with this, it is actually important to have higher density around those regions, which is completely opposite to your claim.

You are right. This sentence ~~was~~ changed in the new version of our paper. With circular influence zones, adapting ~~the size of such zones reduce it~~ in each direction around the record position regardless of the illumination changes. In addition, we think that the oversampling issue is due to the circular shape of the record's influence zone as explained in the figure 2,

* `The accuracy parameter \alpha is set to 1 in equation 11...'
I failed to understand your reasoning. Why is it fine to fix the accuracy parameter?

\alpha is set to 1 because near objects detection is more accurate with our method: axes lengths are

~~initialize~~ by the distance to the closest object in each axis direction.

= Section 5 =
* The visibility term in Equation 12 needs to be clarified as I mentioned above.
* 'S then the radiance is uniform'
Does it mean reflected radiance from S is uniform since we only consider Lambertian?
* '(radiance is not uniform)'
Point light sources have infinite radiance. If you are talking about irradiance due to a point light source, its directionally uniform.
* Is E(R) irradiance at the record position? If so, what is x? Overall, how you actually compute gradients are not very clear from the current explanation in 5.2. Maybe adding a summary of each step would be helpful.

These points are clarified in the new version of our paper.

* 'lighting change is linear'
Equation 17 is clearly not linear. Although Equation 18 uses linear interpolation, it does not result in a linear function because of Equation 17.

On each axis, the interpolation is not linear. However, the irradiance at a point p (given by the equation 18), located between two axes k_1 and k_2, is evaluated using a linear interpolation between 2 not linear interpolated irradiance values E_tilde_k1 and E_tilde_k2.
In the paper, we complete the sentence "We assume that the lighting change is linear" with "between 2 successive axes k_1 and k_2."

= Section 6 =
* 'The actual values for these 2 points are'
What do 'actual values' mean? Are those values used for final rendering?

We have clarified this point in the new version of our paper.

= Section 7 =
* 'When applied to a record R... at R through interpolation'
I did not understand this part. This part should be rewritten in my opinion.

We removed this part in the new version of the paper. The Krivanek's adaptive caching method was briefly explained in the section "Related Work".

* The result with \rho = 1% in Figure 13 still has some artifacts, so it did not really solve the problem.

There are still some artifacts in the figure 13. On the yellow tore, the artifacts are due to to bugs in our lighting simulation software (see the answer to reviewer 1 about artifacts visible in video). On the ground, the shadows quality is good and the small persistent artifacts can be removed with a more restrictive rho parameter.

== Clarity and quality of illustrations: Average ==
Figure 1 is almost irrelevant to the proposed method. Figure 2 and 3 essentially tell the same message to me, so I would just keep Figure 3.

We think that the figure 1 helps to understand how the equivalent point light sources are determined with a Monte Carlo ray tracing. Concerning the figures 1 and 3, they ~~not tell~~ exactly the same message. The figure 2 illustrates the main idea of using non-circular influence zones. The figure 3 shows a comparison between a circular influence zone and our adaptive influence zone constructed with 8 axes around the record position.